SUPERCOMPUTING: IS THE U.S. ON THE RIGHT PATH?

HEARING

BEFORE THE

COMMITTEE ON SCIENCE
HOUSE OF REPRESENTATIVES

ONE HUNDRED EIGHTH CONGRESS

FIRST SESSION

JULY 16, 2003

Serial No. 108–21

Printed for the use of the Committee on Science

Available via the World Wide Web: http://www.house.gov/science

U.S. GOVERNMENT PRINTING OFFICE

88–231PS WASHINGTON : 2003

For sale by the Superintendent of Documents, U.S. Government Printing Office
Internet: bookstore.gpo.gov Phone: toll free (866) 512–1800; DC area (202) 512–1800
Fax: (202) 512–2250 Mail: Stop SSOP, Washington, DC 20402–0001

COMMITTEE ON SCIENCE

HON. SHERWOOD L. BOEHLERT, New York, *Chairman*

LAMAR S. SMITH, Texas
CURT WELDON, Pennsylvania
DANA ROHRABACHER, California
JOE BARTON, Texas
KEN CALVERT, California
NICK SMITH, Michigan
ROSCOE G. BARTLETT, Maryland
VERNON J. EHLERS, Michigan
GIL GUTKNECHT, Minnesota
GEORGE R. NETHERCUTT, JR., Washington
FRANK D. LUCAS, Oklahoma
JUDY BIGGERT, Illinois
WAYNE T. GILCHREST, Maryland
W. TODD AKIN, Missouri
TIMOTHY V. JOHNSON, Illinois
MELISSA A. HART, Pennsylvania
JOHN SULLIVAN, Oklahoma
J. RANDY FORBES, Virginia
PHIL GINGREY, Georgia
ROB BISHOP, Utah
MICHAEL C. BURGESS, Texas
JO BONNER, Alabama
TOM FEENEY, Florida
RANDY NEUGEBAUER, Texas

RALPH M. HALL, Texas
BART GORDON, Tennessee
JERRY F. COSTELLO, Illinois
EDDIE BERNICE JOHNSON, Texas
LYNN C. WOOLSEY, California
NICK LAMPSON, Texas
JOHN B. LARSON, Connecticut
MARK UDALL, Colorado
DAVID WU, Oregon
MICHAEL M. HONDA, California
CHRIS BELL, Texas
BRAD MILLER, North Carolina
LINCOLN DAVIS, Tennessee
SHEILA JACKSON LEE, Texas
ZOE LOFGREN, California
BRAD SHERMAN, California
BRIAN BAIRD, Washington
DENNIS MOORE, Kansas
ANTHONY D. WEINER, New York
JIM MATHESON, Utah
DENNIS A. CARDOZA, California
VACANCY

CONTENTS

July 16, 2003

Page

Appendix 2: Additional Material for the Record

SUPERCOMPUTING: IS THE U.S. ON THE RIGHT PATH?

WEDNESDAY, JULY 16, 2003

HOUSE OF REPRESENTATIVES,
COMMITTEE ON SCIENCE,
Washington, DC.

The Committee met, pursuant to call, at 10:21 a.m., in Room 2318 of the Rayburn House Office Building, Hon. Sherwood L. Boehlert (Chairman of the Committee) presiding.

COMMITTEE ON SCIENCE
U.S. HOUSE OF REPRESENTATIVES

Supercomputing: Is the U.S. on the Right Path?

Wednesday, July 16, 2003
10:00 a.m. - 12:00 Noon
2318 Rayburn House Office Building

Witness List

Dr. Raymond L. Orbach
Director, Office of Science
Department of Energy

Dr. Peter A. Freeman
Assistant Director, Computer and Information Science and Engineering Directorate (CISE)
National Science Foundation (NSF)

Dr. Daniel A. Reed
Director, National Center for Supercomputing Applications (NCSA)
University of Illinois at Urbana-Champaign

Mr. Vincent Scarafino
Manager, Numerically Intensive Computing
Ford Motor Company

HEARING CHARTER

COMMITTEE ON SCIENCE
U.S. HOUSE OF REPRESENTATIVES

Supercomputing: Is the U.S. on the Right Path?

WEDNESDAY, JULY 16, 2003
10:00 A.M.–12:00 P.M.
2318 RAYBURN HOUSE OFFICE BUILDING

1. Purpose

On Wednesday, July 16, 2003, the House Science Committee will hold a hearing to examine whether the United States is losing ground to foreign competitors in the production and use of supercomputers[1] and whether federal agencies' proposed paths for advancing our supercomputing capabilities are adequate to maintain or regain the U.S. lead.

2. Witnesses

Dr. Raymond L. Orbach is the Director of the Office of Science at the Department of Energy. Prior to joining the Department, Dr. Orbach was Chancellor of the University of California at Riverside.

Dr. Peter A. Freeman is Assistant Director for the Computer and Information Science and Engineering Directorate (CISE) at the National Science Foundation (NSF). Prior to joining NSF in 2002, he was professor and founding Dean of the College of Computing at Georgia Institute of Technology.

Dr. Daniel A. Reed is the Director of the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign. NCSA is the leader of one of NSF's two university-based centers for high-performance computing. Dr. Reed is also the Director of the National Computational Science Alliance and is a principal investigator in the National Science Foundation's TeraGrid project. Earlier this year, Dr. Reed was appointed to the President's Information Technology Advisory Committee (PITAC).

Mr. Vincent Scarafino is the Manager of Numerically Intensive Computing at Ford Motor Company, where he focuses on providing flexible and reliable supercomputer resources for Ford's vehicle product development, including vehicle design and safety analysis.

3. Overarching Questions

The hearing will address the following overarching questions:

1. Is the U.S. losing its leadership position in supercomputing? Do the available supercomputers allow United States science and industry to be competitive internationally? Are federal efforts appropriately targeted to deal with this challenge?
2. Are federal agencies pursuing conflicting supercomputing programs? What can be done to ensure that federal agencies pursue a coordinated policy for providing supercomputing to meet the future needs for science, industry, and national defense?
3. Is the National Science Foundation moving away from the policies and programs that in the past have provided broad national access to advanced supercomputers?
4. Can the U.S. fulfill its scientific and defense supercomputing needs if it continues to rely on machines designed for mass-market commercial applications?

[1] Supercomputing is also referred to as high-performance computing, high-end computing, and sometimes advanced scientific computing.

4. Brief Overview

- High-performance computers (also called supercomputers) are an essential component of U.S. scientific, industrial, and military competitiveness. However, the fastest and most efficient supercomputer in the world today is in Japan, not the U.S. Some experts claim that Japan was able to produce a computer so far ahead of the American machines because the U.S. had taken an overly cautious or conventional approach for developing new high-performance computing capabilities.
- Users of high-performance computing are spread throughout government, industry, and academia, and different high-performance computing applications are better suited to different types of machines. As the U.S. works to develop new high-performance computing capabilities, extraordinary coordination among agencies and between government and industry will be required to ensure that creative new capabilities are developed efficiently and that all of the scientific, governmental, and industrial users have access to the high-performance computing hardware and software best suited to their applications.
- The National Science Foundation (NSF) currently provides support for three supercomputing centers: the San Diego Supercomputer Center, the National Center for Supercomputing Applications at Urbana-Champaign in Illinois, and the Pittsburgh Supercomputing Center. These centers, along with their partners at other universities, are the primary source of high-performance computing for researchers in many fields of science. Currently, support for these centers beyond fiscal year 2004 is uncertain, and in the past few years NSF has been increasing its investment in a nationwide computing grid, in which fast connections are built between many computers to allow for certain types of high-performance scientific computing and advanced communications and data management. It is not clear whether this "grid computing" approach will provide the high-performance computing capabilities needed in all the scientific fields that currently rely on the NSF supercomputing centers.
- At the Department of Energy, there are two programs aimed at advancing high-performance computing capabilities. One, in the National Nuclear Security Administration (NNSA), is the continuation of a long-term effort to provide supercomputers to be used for modeling nuclear weapons effects; these simulations are particularly important in light of existing bans on nuclear weapon testing. In the other program, the Office of Science is now proposing to supplement its current advanced scientific computing activities with a new effort designed to create the world's fastest supercomputers.

5. Current Issues

Is the U.S. Competitive?

Japan's Earth Simulator is designed to perform simulations of the global environment that allow researchers to study scientific questions related to climate, weather, and earthquakes. It was built by NEC for the Japanese government at a cost of at least $350 million and has been the fastest computer in the world since it began running in March 2002. When the first measures of its speed were performed in April 2002, researchers determined that the Earth Simulator was almost five times faster than the former record holder, the ASCI White System at Lawrence Livermore National Laboratory, and also used the machine's computing power significantly more efficiently.[2]

This new development caused a great deal of soul-searching in the high-performance computing community about the U.S. approach to developing new capabilities and the emphasis on using commercially available (not specialized or custom-made) components. Going forward, it is not clear whether or not such a commodity-based approach will allow the U.S. high-performance computing industry to remain competitive. It is also unclear if the new machines produced by this approach will be able provide American academic, industrial, and governmental users with the high-performance computing capabilities they need to remain the best in the world in all critical applications.

[2] For the U.S. supercomputers, typical scientific applications usually only are able to utilize 5–10 percent of the theoretical maximum computing power, while the design of the Earth Simulator makes 30–50 percent of its power accessible to the majority of typical scientific applications.

Will All Users Be Served?

Users of high-performance computing are spread throughout government, industry, and academia. Different high-performance computing applications are better suited to different types of machines. For example, weather modeling and simulations of nuclear weapons require many closely-related calculations, so machines for these applications must have components that communicate with each other quickly and often. Other applications, such as simulations of how proteins fold, can be efficiently performed with a more distributed approach on machines in which each component tackles a small piece of the problem and works in relative isolation. In the U.S., the major producers of high-performance computers include IBM, Hewlett-Packard, and Silicon Graphics, Inc., whose products lean toward the more distributed approach, and Cray, whose products are more suited to problems that require the performance of closely-related calculations. The Japanese (NEC, Fujitsu, and Hitachi), also produce this sort of machine. The concern is that the U.S. on the whole has moved away from developing and manufacturing the machines needed for problems with closely-related calculations because the more distributed machines have a bigger commercial market. The Japanese have been filling this gap, but the gap could still impact the access of American scientists to the types of supercomputers that they need for certain important research problems.

Responsibility for providing high-performance computing capabilities to existing users and for developing new capabilities is distributed among 11 different federal agencies and offices and relies heavily on industry for development and production. In this environment, extraordinary amounts of coordination are needed to ensure that new capabilities are developed efficiently and that the most appropriate kinds of hardware and software are available to the relevant users—coordination among agencies and between government and industry, as well as cooperation among universities and hardware and software companies. The results of an ongoing interagency effort to produce a coherent high-performance computing roadmap and the influence this roadmap has on agencies' programs will be the first test.

Where are the DOE Office of Science and the NSF Programs Headed?

Both NSF and the DOE Office of Science are moving ahead in significant new directions. At NSF, no plans have been announced to continue the Partnerships for Advanced Computational Infrastructure program, which supports the supercomputer centers, beyond fiscal year 2004. In addition, a proposed reorganization of NSF's Computer and Information Sciences and Engineering Directorate was announced on July 9 that includes a merging of the Advanced Computational Infrastructure program (which includes the support for the supercomputer centers) and the Advanced Networking Infrastructure program (which supports efforts on grid computing—an alternative approach to high-performance computing). Some scientists have expressed concerns that NSF may be reducing its commitment to providing researchers with a broad range of supercomputing capabilities and instead focusing its attention on grid computing and other distributed approaches.

For the DOE Office of Science, the fiscal year 2004 budget request proposes a new effort in next-generation computer architecture to identify and address major bottlenecks in the performance of existing and planned DOE science applications. In addition, the July 8 mark-up of the House Energy and Water Development Appropriations Subcommittee sets funding for the Advanced Scientific Computing Research initiative at \$213.5 million, an increase of \$40 million over the request and \$46 million over the previous year. Decisions about the future directions for high-performance computing at NSF and DOE Office of Science are clearly being made now.

The White House has an interagency effort underway, the High End Computing Revitalization Task Force (HECRTF), which is supposed to result in the agencies' submitting coordinated budget requests in this area for fiscal year 2005.

6. Background

What is High-Performance Computing? High-performance computing—also called supercomputing, high-end computing, and sometimes advanced scientific computing—is a phrase used to describe machines or groups of machines that can perform very complex computations very quickly. These machines are used to solve complicated and challenging scientific and engineering problems or manage large amounts of data. There is no set definition of how fast a computer must be to be "high-performance" or "super," as the relevant technologies improve so quickly that the high-performance computing achievements of a few years ago could be handled now by today's desktops. Currently, the fastest supercomputers are able to perform trillions of calculations per second.

What is High-Performance Computing Used For? High-performance computing is needed for a variety of scientific, industrial, and national defense applications. Most often, these machines are used to simulate a physical system that is difficult to study experimentally. The goal can be to use the simulation as an alternative to actual experiments (e.g., for nuclear weapon testing and climate modeling), as a way to test our understanding of a system (e.g., for particle physics and astrophysics), or as a way to increase the efficiency of future experiments or product design processes (e.g., for development of new industrial materials or fusion reactors). Other major uses for supercomputers include performing massive complex mathematical calculations (e.g., for cryptanalysis) or managing massive amounts of data (e.g., for government personnel databases).

Scientific Applications: There are a rich variety of scientific problems being tackled using high-performance computing. Large-scale climate modeling is used to examine possible causes and future scenarios related to global warming. In biology and biomedical sciences, researchers perform simulations of protein structure, folding, and interaction dynamics and also model blood flows. Astrophysics model planet formation and supernova, and cosmologists analyze data on light from the early universe. Particle physicists use the ultra-fast computers to perform the complex calculations needed to study quantum chromodynamics and improve our understanding of electrons and quarks, the basic building blocks of all matter. Geologists model the stresses within the earth to study plate tectonics, while civil engineers simulate the impact of earthquakes.

National Defense Applications: There are a number of ways in which high-performance computing is used for national defense applications. The National Security Agency (NSA) is a major user and developer of high-performance computers for executing specialized tasks relevant to cryptanalysis (such as factoring large numbers). The Department of Energy's National Nuclear Security Administration is also a major user and developer of machines to be used for designing and modeling nuclear weapons. Other applications within the Department of Defense include armor penetration modeling, weather forecasting, and aerodynamics modeling. Many of the scientific applications also have direct or future defense applications. For example, computational fluid dynamics studies are also of interest to the military, e.g. for modeling turbulence around aircraft. The importance of high-performance computing in many military areas, including nuclear and conventional weapons design, means that machines that alone or when wired together are capable of superior performance at military tasks are subject to U.S. export controls.

Industrial Applications: Companies use high-performance computing in a variety of ways. The automotive industry uses fast machines to maximize the effectiveness of computer-aided design and engineering. Pixar uses massive computer animation programs to produce films. Pharmaceutical companies simulate chemical interactions to help with drug design. The commercial satellite industry needs to manage huge amounts of data for mapping. Financial companies and other industries use large computers to process immense and unpredictable Web transaction volumes, mine databases for sales patterns or fraud, and measure the risk of complex investment portfolios.

What Types of High-Performance Computers Are There? All of the above examples of high-performance computing applications require very fast machines, but they do not all require the same type of very fast machine. There are a number of different ways to build high-performance computers, and different configurations are better suited to different problems. There are many possible configurations, but they can be roughly divided into two classes: big, single-location machines and distributed collections of many computers (this approach is often called grid computing). Each approach has its benefits—the big machines can be designed for a specific problem and are often faster, while grid computing is attractive in part because by using a multitude of commercially-available computers, the purchase and storage cost is often lower than for a large specialized supercomputer.

Since the late 1990's, the U.S. approach to developing new capabilities has emphasized using commercially available (not specialized) components as much as possible. This emphasis has resulted in an increased focus on grid computing, and, in large machines, has led to a hybrid approach in which companies use commercial processors (whose speed is increasing rapidly anyway) to build the machines and then further speed them up by increasing the number of processors and improving the speed at which information is passed between processors. There are a number of distinctions that can be made among large machines bases on how the processors are connected. The differences relate to how fast and how often the various compo-

nents of the computer communicate with each other and how calculations are distributed among the components.

Users thus have a number of options for high-performance computing. Each user must take into account all of the pros and cons of the different configurations when he is deciding what sort of machine to use and how to design software to allow that machine to most efficiently solve his problem. For example, some problems, like weather and climate modeling and cryptanalysis, require lots of communication among computer components and large quantities of stored data, while other applications, like large-scale data analysis for high energy physics experiments or bioinformatics projects, can be more efficiently performed on distributed machines each tackling its own piece of the problem in relative isolation.

How Do Government and Industry Provide Existing and New High-Performance Computing Capabilities? The development and production of high-performance computing capabilities requires significant effort by both government and industry. For any of the applications of high-performance computing described above, the users need good hardware (the high-performance machine or group of machines) and good software (programs that allow them to perform their calculations as accurately and efficiently as possible).

The role of government therefore includes (1) funding research on new approaches to building high-performance computing hardware, (2) in some cases, funding the development stage of that hardware (usually through security agencies), (3) purchasing the hardware to be used by researchers at universities and personnel at government agencies, (4) funding research on software and programs to use existing and new high-performance computing capabilities, and (5) supporting research that actually uses the hardware and software. The role of industry is complementary—i.e., it receives funding to do research and development on new hardware and software, and it is the seller of this hardware and software to government agencies, universities, and companies. The primary industries involved in producing high-performance computing capabilities are computer makers (such as IBM, Hewlett-Packard, Silicon Graphics, Inc., and Cray), chip makers (such as Intel), and software designers. Congress has long had concerns about the health of the U.S. supercomputing industry. In 1996, when the National Center for Atmospheric Research, a privately-run, federally-funded research center, tried to order a supercomputer from NEC for climate modeling, Congress blocked the purchase.

Federal High-Performance Computing Programs: In 1991, Congress passed the High Performance Computing Act, establishing an interagency initiative (now called National Information Technology Research and Development (NITRD) programs) and a National Coordination Office for this effort. Currently 11 agencies or offices participate in the high-end computing elements of the NITRD program (See Table 1 in the appendix). The total requested by all 11 agencies in fiscal year 2003 for high-end computing was $846.5 million. The largest research and development programs are at the National Science Foundation (NSF), which requested $283.5 million, and the Department of Energy Office of Science, which requested $137.8 million. Other major agency activities (all between $80 and $100 million) are at the National Institutes of Health (NIH), the Defense Advanced Research Projects Agency (DARPA), the National Aeronautics and Space Administration (NASA), and the Department of Energy's National Nuclear Security Administration (NNSA). Different agencies concentrate on serving different user communities and on different stages of hardware and software development and application. (In addition to the research and development-type activities that are counted for the data included in Table 1 and referenced above, many agencies, such as NNSA and the National Oceanic and Atmospheric Administration (NOAA), devote significant funding to the purchase and operation of high-performance computers that perform these agencies' mission-critical applications.)[3]

> *National Science Foundation:* The NSF serves a very wide variety of scientific fields within the academic research community, mainly through a series of supercomputing centers, originally established in 1985 and currently funded under the Partnerships for Advanced Computational Infrastructure (PACI) program. The supercomputer centers provide researchers not only with access to high-performance computing capabilities but also with tools and expertise on how best to utilize these resources. The NSF also is supporting the development of the Extensible Terascale Facility (ETF), a nationwide grid of machines that can be used for high-performance computing and advanced communications and

[3] For example, in FY 2003 NOAA spent $36 million on supercomputers—$10 million for machines for climate modeling and $26 million for machines for the National Weather Service.

data management. Recently, some researchers within the high-performance computing community have expressed concern that NSF may be reducing its commitment to the supercomputer centers and increasing its focus on grid computing and distributed approaches to high-performance computing, such as would be used in the ETF.

Department of Energy: The Department of Energy has been a major force in advancing high-performance computing for many years, and the unveiling of the fastest computer in the world in Japan in 2002 resulted in serious self-evaluation at the department, followed by a rededication to efforts to enhance U.S. supercomputing capabilities. The Department of Energy has two separate programs focused on both developing and applying high-performance computing. The Advanced Scientific Computing Research (ASCR) program in the Office of Science funds research in applied mathematics (to develop methods to model complex physical and biological systems), in network and computer sciences, and in advanced computing software tools. For fiscal year 2004, the department has proposed a new program on next-generation architectures for high-performance computing. The Accelerated Strategic Computing Initiative (ASCI) is part of the NNSA's efforts to provide advanced simulation and computing technologies for weapons modeling.

DARPA: DARPA traditionally focuses on the development of new hardware, including research into new architectures and early development of new systems. On July 8, DARPA announced that Cray, IBM, and Sun Microsystems had been selected as the three contractor teams for the second phase of the High Productivity Computing Systems program, in which the goal is to provide a new generation of economically viable, scalable, high productivity computing systems for the national security and industrial user communities in the 2009 to 2010 timeframe.

Other Agencies: NIH, NASA, and NOAA are all primarily users of high performance computing. NIH manages and analyzes biomedical data and models biological processes. NOAA uses simulations to do weather forecasting and climate change modeling. NASA has a variety of applications, including atmospheric modeling, aerodynamic simulations, and data analysis and visualization. The National Security Agency (NSA) both develops and uses high-performance computing for a number of applications, including cryptanalysis. As a user, NSA has a significant impact on the high-performance computing market, but due to the classified nature of its work, the size of its contributions to High End Computing Infrastructure and Applications and the amount of funding it uses for actual operation of computers is not included in any of the data.

Interagency Coordination: The National Coordination Office (NCO) coordinates planning, budget, and assessment activities for the Federal Networking and NITRD Program through a number of interagency working groups. The NCO reports to the White House Office of Science and Technology Policy and the National Science and Technology Council. In 2003, NCO is also managing the High End Computing Revitalization Task Force (HECRTF), an interagency effort on the future of U.S. high-performance computing. The HECRTF is tasked with development of a roadmap for the interagency research and development for high-end computing core technologies, a federal high-end computing capacity and accessibility improvement plan, and a discussion of issues relating to federal procurement of high-end computing systems. The product of the HECRTF process is expected to guide future investments in this area, starting with agency budget submissions for fiscal year 2005.

The Role of Industry: Industry plays a critical role in developing and providing high-performance computing capabilities to scientific, industrial, and defense users. Many supercomputers are purchased directly from computer companies like IBM, Hewlett-Packard, Silicon Graphics, Inc., and Cray, and the groups that do build their own high-performance clusters do so from commercially available computers and workstations. Industry is a recipient of federal funding for initial research into new architectures for hardware, for development of new machines, and for production of standard and customized systems for government and universities, but industry also devotes its own funding to support research and development. The research programs do not just benefit the high-performance computing community, as new architectures and faster chips lay the groundwork for better performing computers and processors in all commercial information technology products.

The State of the Art in High-Performance Computing: Twice a year, a list of the 500 fastest supercomputers is compiled; the latest list was released on June 23, 2003

(see Table 2 in the appendix).[4] The Earth Simulator supercomputer, built by NEC and installed last year at the Earth Simulator Center in Yokohama, Japan, continues to hold the top spot as the best performer. It is approximately twice as fast as the second place machine, the ASCI Q system at Los Alamos National Laboratory, built by Hewlett-Packard. Of the top twenty machines, eight are located at various Department of Energy national laboratories and two at U.S. universities,[5] and nine were made by IBM and five by Hewlett-Packard.

7. Witness Questions

The witnesses were asked to address the following questions in their testimony:

Questions for Dr. Raymond L. Orbach

- The Office of Science appears to have embarked on a new effort in next-generation advanced scientific computer architecture that differs from the development path currently pursued by the National Nuclear Security Agency (NNSA), the lead developer for advanced computational capability at the Department of Energy (DOE). Why is the Office of Science taking this approach?
- How is the Office of Science cooperating with the Defense Advanced Research Projects Agency, which supports the development of advanced computers for use by the National Security Agency and other agencies within the Department of Defense?
- To what extent will the Office of Science be guided by the recommendations of the High-End Computing Revitalization Task Force? How will the Office of Science contribute to the Office of Science and Technology Policy plan to revitalize high-end computing?
- To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any divergence on the advanced computing development programs at the Office of Science?

Questions for Dr. Peter A. Freeman

- Some researchers within the computer science community have suggested that the NSF may be reducing its commitment to the supercomputer centers. Is this the case? To what extent does the focus on grid computing represent a move away from providing researchers with access to the most advanced computing equipment?
- What are the National Science Foundation's (NSF's) plans for funding the supercomputer centers beyond fiscal year 2004? To what extent will you be guided by the recommendation of the NSF Advisory Panel on Cyberinfrastructure to maintain the Partnerships for Advanced Computational Infrastructure, which currently support the supercomputer centers?
- To what extent will NSF be guided by the recommendations of the High-End Computing Revitalization Task Force? How will NSF contribute to the Office of Science and Technology Policy plan to revitalize high-end computing?
- To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any such divergence on the advanced computing programs at NSF?

Questions for Dr. Daniel A. Reed

- Some researchers within the computer science community have suggested that the National Science Foundation (NSF) may be reducing its commitment to provide advanced scientific computational capability to U.S. scientists and engineers. Have you detected any change in policy on the part of NSF?
- What advanced computing capabilities must the Federal Government provide the academic research community for the government's programs to be con-

[4] The top 500 list is compiled by researchers at the University of Mannheim (Germany), Lawrence Berkeley National Laboratory, and the University of Tennessee and is available on line at *http://www.top500.org/*. For a machine to be included on this public list, its owners must send information about its configuration and performance to the list-keepers. Therefore, the list is not an entirely comprehensive picture of the high-performance computing world, as classified machines, such as those used by NSA, are not included.

[5] The two university machines are located at the Pittsburgh Supercomputing Center (supported primarily by NSF) and Louisiana State University's Center for Applied Information Technology and Learning. The remaining 12 machines include four in Europe, two in Japan, and one each at the National Oceanic & Atmospheric Administration, the National Center for Atmospheric Research, the Naval Oceanographic Office, and NASA.

sidered successful? Are the programs for developing the next-generation of advanced scientific computing that are currently underway at government agencies on track to provide these capabilities? If not, why not?

- For academic scientists and engineers, what is the difference between the advanced scientific computing capabilities provided by NSF and those provided by the Department of Energy?

Questions for Mr. Vincent F. Scarafino

- How does Ford use high-performance computing? How do computing capabilities affect Ford's competitiveness nationally and internationally?
- What does Ford see as the role of the Federal Government in advancing high-performance computing capabilities and in making these capabilities accessible to users? Are current agency programs for developing the next-generation of advanced scientific computing adequate to provide these capabilities? If not, why not?
- Is the U.S. government cooperating appropriately with the private sector on high-performance computing, and is the level of cooperation adequate to sustain leadership and meet scientific and industrial needs?
- To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any divergence on Ford's access to advanced computing capabilities?

APPENDIX

Table 1a: Fiscal Year 2003 Budget Requests for High End Computing by Agencies Participating in the National Information Technology Research and Development program (dollars in millions)

Agency	High End Computing: Infrastructure and Applications	High End Computing: Research and Development	Total for High End Computing
DARPA	16.8	81.9	98.7
DOE/NNSA	41.4	39.5	80.9
DOE Office of Science	**98.5**	**39.3**	**137.8**
EPA	1.8	0.0	1.8
NASA	68.4	26.0	94.4
NIH	88.2	8.9	97.1
NIST	3.5	0.0	3.5
NOAA	13.3	1.8	15.1
NSA	--	31.9	31.9
NSF	**215.2**	**68.3**	**283.5**
ODDR&E	--	1.8	1.8
Total:	**547.1**	**299.4**	**846.5**

Source: NITRD National Coordination Office Fiscal Year 2003 Blue Book. The Blue Book is released in August of each year, and thus the data on FY 2003 spending and FY 2004 budget requests levels has not yet been provided to the National Coordination Office.

Note: In addition to the research and development-type activities that are counted for the data included in this table and Table 1b, many agencies devote significant funding to the purchase and operation of high-performance computers that perform these agencies' mission-critical applications.

Acronyms: DARPA—Defense Advanced Research Projects Agency, DOE/NNSA—Department of Energy's National Nuclear Security Administration, EPA—Environmental Protection Agency, NASA—National Aeronautics and Space Administration, NIH—National Institutes of Health, NIST—National Institute of Standards and Technology, NOAA—National Oceanic and Atmospheric Administration, NSA—National Security Agency, NSF—National Science Foundation, ODDR&E—Office of the Director of Defense Research and Engineering, VA—Department of Veterans Affairs.

Table 1b: Funding History from fiscal year 1992 to fiscal year 2003 of high-performance computing research and development programs at various agencies.

	FY 1992	FY 1993	FY 1994	FY 1995	FY 1996	FY 1997	FY 1998	FY 1999	FY 2000	FY 2001	FY 2002	FY 2003 (Requests)
DARPA	141.80	169.20	136.20	142.70	77.96	72.70	84.80	48.00	36.50	96.20	81.30	98.70
DOE/NNSA									113.90	168.30	75.60	80.90
DOE/SC	**73.00**	**76.20**	**84.60**	**73.10**	**84.49**	**86.00**	**90.53**	**91.90**	**84.10**	**130.30**	**126.70**	**137.80**
EPA	4.50	6.10	5.90	10.50	8.70	5.60	5.38	4.20	3.90	3.50	1.80	1.80
NASA	64.00	70.20	84.60	87.40	75.55	88.00	90.10	71.40	124.80	86.80	62.10	94.40
NIH	8.90	34.40	29.50	29.90	22.40	23.40	23.74	27.10	34.10	59.50	87.20	97.10
NIST	0.90	0.90	0.90	3.60	5.59	4.00	3.99	3.50	3.50	3.50	3.50	3.50
NOAA	1.80	9.40	9.80	2.80	3.30	4.30	4.30	8.80	13.20	12.00	15.60	15.10
NSA		40.20	32.70	28.20	29.48	30.40	26.42	24.00	31.70	32.90	41.60	31.90
NSF	**127.00**	**133.90**	**139.10**	**150.00**	**140.32**	**129.20**	**132.90**	**224.70**	**289.80**	**311.70**	**291.50**	**283.50**
ODDR&E									2.00	2.00	2.00	1.80
VA					3.00	1.00						
Totals	**421.90**	**540.50**	**523.30**	**528.20**	**450.79**	**444.60**	**462.16**	**503.60**	**737.50**	**906.70**	**788.90**	**846.50**

Source: NITRD National Coordination Office Blue Books, Fiscal Years 1992 to 2003.

Acronyms: DARPA—Defense Advanced Research Projects Agency, DOE/NNSA—Department of Energy's National Nuclear Security Administration, DOE/SC—Department of Energy's Office of Science, EPA—Environmental Protection Agency, NASA—National Aeronautics and Space Administration, NIH—National Institutes of Health, NIST—National Institute of Standards and Technology, NOAA—National Oceanic and Atmospheric Administration, NSA—National Security Agency, NSF—National Science Foundation, ODDR&E—Office of the Director of Defense Research and Engineering, VA—Department of Veterans Affairs.

Program History: Figures from FY 1992-1995 reflect the funding for the High Performance Computing Systems and the Advanced Software Technology and Algorithms Programs. Figures from FY 1996-1999 reflect the funding for the High End Computing and Computation Program. Figures from FY 2000-2003 reflect the funding for the High End Computing Infrastructure and Applications and Research and Development Programs.

Table 2: The top twenty machines of the TOP500 List of the World's Fastest Supercomputers (full list available on line at http://www.top500.org/).

Rank	Manufacturer Computer/Number of Processors	Installation Site Country/Year
1	**NEC** Earth-Simulator/5120	**Earth Simulator Center** Japan/2002
2	**Hewlett-Packard** ASCI Q - AlphaServer SC ES45/1.25 GHz/8192	**Los Alamos National Laboratory** Los Alamos, NM, USA/2002
3	**Linux Networx** MCR Linux Cluster Xeon 2.4 GHz-Quadrics/2304	**Lawrence Livermore National Laboratory** Livermore, CA, USA/2002
4	**IBM** ASCI White, SP Power3 375 MHz/8192	**Lawrence Livermore National Laboratory** Livermore, CA, USA/2000
5	**IBM** SP Power3 375 MHz 16 way/6656	**NERSC/Lawrence Berkeley National Laboratory** Berkeley, CA, USA/2002
6	**IBM** xSeries Cluster Xeon 2.4 GHz - Quadrics/1920	**Lawrence Livermore National Laboratory** Livermore, CA, USA/2003
7	**Fujitsu** PRIMEPOWER HPC2500 (1.3 GHz)/2304	**National Aerospace Laboratory of Japan** Japan/2002
8	**Hewlett-Packard** rx2600 Itanium2 1 GHz Cluster - Quadrics/1540	**Pacific Northwest National Laboratory** Richland, WA, USA/2003
9	**Hewlett-Packard** AlphaServer SC ES45/1 GHz/3016	**Pittsburgh Supercomputing Center** Pittsburgh, PA, USA/2001
10	**Hewlett-Packard** AlphaServer SC ES45/1 GHz/2560	**Commissariat a l'Energie Atomique (CEA)** France/2001
11	**HPTi** Aspen Systems, Dual Xeon 2.2GHz-Myrinet2000/1536	**Forecast Systems Laboratory - NOAA** Boulder, CO, USA/2002
12	**IBM** pSeries 690 Turbo 1.3GHz/1280	**HPCx (UK Academic Research Center)** UK/2002
13	**IBM** pSeries 690 Turbo 1.3GHz/1216	**NCAR (National Center for Atmospheric Research)** Boulder, CO, USA/2002
14	**IBM** pSeries 690 Turbo 1.3GHz/1184	**Naval Oceanographic Office (NAVOCEANO)** Stennis, MS, USA/2002
15	**IBM** pSeries 690 Turbo 1.3GHz/960	**European Ctr. for Medium-Range Weather Forecasts** UK/2002
16	**IBM** pSeries 690 Turbo 1.3GHz/960	**European Ctr. for Medium-Range Weather Forecasts** UK/2002
17	**Intel** ASCI Red/9632	**Sandia National Laboratories** Albuquerque, NM, USA/1999
18	**IBM** pSeries 690 Turbo 1.3GHz/864	**Oak Ridge National Laboratory** Oak Ridge, TN, USA/2002
19	**Atipa Technology** P4 Xeon 1.8 GHz - Myrinet/1024	**Louisiana State University** Baton Rouge, LA, USA/2002
20	**Hewlett-Packard** AlphaServer SC ES45/1 GHz/1392	**NASA/Goddard Space Flight Center** Greenbelt, MD, USA/2002

Chairman BOEHLERT. The hearing will come to order. And I apologize for being the delinquent Member of the group gathered here. My apologies to all.

It is a pleasure to welcome everyone here this morning for this important hearing. At first blush, today's topic, supercomputing, may seem technical and arcane, of interest to just a few researchers who spend their lives in the most rarefied fields of science. But in reality, the subject of this hearing is simple and accessible, and it has an impact on all of us, because supercomputing affects the American economy and our daily lives, perhaps more so than so many, many other things that we focus a lot of time and attention on.

Supercomputers help design our cars, predict our weather, and deepen our understanding of the natural forces that govern our lives, such as our climate. Indeed, computation is now widely viewed as a third way of doing science; building on the traditional areas of theory and experimentation.

So when we hear that the U.S. may be losing its lead in supercomputing, that Japan now has the fastest supercomputer, that the U.S. may be returning to a time when our top scientists didn't have access to the best machines, that our government may have too fragmented a supercomputing policy, well, those are issues a red flag should be waved on to capture the attention of all of us.

And those issues have captured our attention. The purpose of this hearing is to gauge the state of U.S. supercomputing and to determine how to deal with any emerging problems.

I don't want to exaggerate, we are not at a point of crisis. Most of the world's supercomputers are still made by and used by Americans, but we are at a pivotal point when we need to make critical decisions to make sure that remains the case.

And maintaining U.S. leadership requires a coordinated, concerted effort by the Federal Government. Let me stress that. Coordinated, concerted effort by the Federal Government. The Federal Government has long underwritten the basic research that fuels the computer industry, has purchased the highest end computers, and has ensured that those computers are available to a wide range of American researchers. This Committee has played an especially crucial role in ensuring access, pushing for the creation of the National Science Foundation Supercomputer Centers back in the early '80's.

Government action is just as needed now. But what action? The Department of Energy is proposing to move away from our reliance on more mass-market supercomputers to pursue research on massive machines designed to solve especially complex problems. NSF appears to be moving away from super—supporting supercomputer centers to a more distributed computing approach. These policies need to be examined.

So with that in mind, here are some of the questions we intend to pursue today. Is the U.S. losing its lead in supercomputing, and what can be done about it? What federal policies should be pursued to maintain our lead, and how should we judge whether they are succeeding? Is federal policy sufficiently coordinated? And I think the answer is clear. And are the new directions being pursued by

NSF and the Department of Energy the proper approach to maintaining our lead?

We have a distinguished group of experts, and I look forward to hearing their testimony.

With that, it is my pleasure to yield to the distinguished Ranking Member, Mr. Hall of Texas.

[The prepared statement of Mr. Boehlert follows:]

PREPARED STATEMENT OF CHAIRMAN SHERWOOD L. BOEHLERT

It's a pleasure to welcome everyone here this morning for this important hearing. At first blush, today's topic, supercomputing, may seem technical and arcane—of interest to just a few researchers who spend their lives in the most rarefied fields of science. But in reality, the subject of this hearing is simple and accessible, and it has an impact on all of us because supercomputing affects the American economy and our daily lives.

Supercomputers help design our cars, predict our weather, and deepen our understanding of the natural forces that govern our lives, such as our climate. Indeed, computation is now widely viewed as a third way of doing science—building on the traditional areas of theory and experimentation.

So when we hear that the U.S. may be losing its lead in supercomputing, that Japan now has the fastest supercomputer, that the U.S. may be returning to a time when our top scientists didn't have access to the best machines, that our government may have too fragmented a supercomputing policy—well, those issues are a red flag that should capture the attention of all of us.

And those issues have captured our attention. The purpose of this hearing is to gauge the state of U.S. supercomputing and to determine how to deal with any emerging problems.

I don't want to exaggerate—we're not at a point of crisis—most of the world's supercomputers are still made by, and used by Americans. But we are at a pivotal point when we need to make critical decisions to make sure that remains the case.

And maintaining U.S. leadership requires a coordinated, concerted effort by the Federal Government. The Federal Government has long underwritten the basic research that fuels the computer industry, has purchased the highest-end computers, and has ensured that those computers are available to a wide range of American researchers. This committee has played an especially crucial role in ensuring access, pushing for the creation of the National Science Foundation (NSF) Supercomputer Centers back in the early '80s.

Government action is just as needed now. But what action? The Department of Energy is proposing to move away from our reliance on more mass-market supercomputers to pursue research on massive machines design to solve especially complex problems. NSF appears to be moving away from supporting supercomputer centers to a more distributed computing approach. These policies need to be examined.

So, with that in mind, here are some of the questions we intend to pursue today:

- Is the U.S. losing its lead in supercomputing and what can be done about that?
- What federal policies should be pursued to maintain our lead and how should we judge whether they are succeeding?
- Is federal policy sufficiently coordinated and are the new directions being pursued by NSF and the Department of Energy the proper approach to maintaining our lead?

We have a distinguished group of experts, and I look forward to hearing their testimony.

Mr. HALL. Mr. Chairman, thank you. And I am pleased to join you today in welcoming our witnesses. And I thank you for your time, not only in your appearing here, but in preparation and travel. And thank you for your usual courtesy.

Computation has become one of the, I guess, principal tools, along with the theory and experiment for conducting science and engineering research and development. There is no question that the U.S. preeminence in science and technology will not and cannot

continue unless our scientists and engineers have access to the most powerful computers available.

The Science Committee has had a deep and sustained interest in this subject since the emergence of supercomputing in the late 1970's. And the initial concern of the Committee was to ensure that the U.S. scientists and engineers, outside of the classified research world, had access to the most powerful computers. We have supported programs to provide this access, such as the supercomputer centers program at NSF.

Moreover, the Committee has encouraged the efforts of the Federal R&D agencies to develop a coordinated R&D program to accelerate computing and networking developments. The High Performance Computing Act of 1991 formalized this interagency R&D planning and coordination process.

The value and importance of the resulting interagency information technology R&D program is quite evident from its appearances as a formal presidential budget initiative through three different presidential administrations.

So today, I think we want to assess a particular component of the federal information technology R&D effort. That is, the pathway being followed for the development of high-end computers and the provision being made to provide access to these machines by the U.S. research community.

Questions have been raised as to whether we put all of our eggs in one basket by mainly focusing on supercomputers based on commodity components. The recent success of the specialized Japanese Earth Simulator computer has triggered a review of the computing needs of the scientific and technical community and reconsideration of the R&D and acquisition plan needed for the next several years and for the longer-term.

So we would be very interested today in hearing from our witnesses about where we are now in terms of high-end computing capabilities and where we should be going to provide the kinds of computer systems needed to tackle the most important and certainly challenging of all problems.

I also want to explore what the roles of the various federal agencies ought to be in the development of new classes of high-end computers and for providing access for the general U.S. research community to these essential tools.

I appreciate the attendance of our witnesses, and I look forward to your discussion.

I yield back my time.

[The prepared statement of Mr. Hall follows:]

PREPARED STATEMENT OF REPRESENTATIVE RALPH M. HALL

Mr. Chairman, I am pleased to join you today in welcoming our witnesses, and I congratulate you on calling this hearing on federal R&D in support of high-performance computing.

Computation has become one of the principal tools, along with theory and experiment, for conducting science and engineering research and development. There is no question that U.S. preeminence in science and technology will not continue unless our scientists and engineers have access to the most powerful computers available.

The Science Committee has had a deep and sustained interest in this subject since the emergence of supercomputing in the late 1970s.

The initial concern of the Committee was to ensure that U.S. scientists and engineers, outside of the classified research world, had access to the most powerful computers. we have supported programs to provide this access, such as the supercomputer centers program at NSF.

Moreover, the Committee has encouraged the efforts of the federal R&D agencies to develop a coordinated R&D program to accelerate computing and networking developments. The High Performance Computing Act of 1991 formalized this interagency R&D planning and coordination process.

The value and importance of the resulting interagency information technology R&D program is evident from its appearance as a formal presidential budget initiative through 3 different Administrations.

Today, we want to assess a particular component of the federal information technology R&D effort. That is, the pathway being followed for the development of high-end computers and the provisions being made to provide access to these machines by the U.S. research community.

Questions have been raised as to whether we have put all of our eggs into one basket by mainly focusing on supercomputers based on commodity components. The recent success of the specialized Japanese Earth Simulator computer has triggered a review of the computing needs of the scientific and technical community and a reconsideration of the R&D and acquisition plan needed for the next several years, and for the longer-term.

I will be interested in hearing from our witnesses about where we are now in terms of high-end computing capabilities and where we should be going to provide the kinds of computer systems needed to tackle the most important and computationally challenging problems.

I also want to explore what the roles of the various federal agencies ought to be in the development of new classes of high-end computers and for providing access for the general U.S. research community to these essential tools.

I appreciate the attendance of our witnesses, and I look forward to our discussion.

Chairman BOEHLERT. Thank you very much, Mr. Hall. And without objection, all other's opening statements will be made a part of the record at this juncture.

[The prepared statement of Ms. Johnson follows:]

PREPARED STATEMENT OF REPRESENTATIVE EDDIE BERNICE JOHNSON

Thank you, Chairman for calling hearing to examine the very important issue of Supercomputing. I also want to thank our witnesses for agreeing to appear today.

We are here to discuss whether the United States is losing ground to foreign competitors in the production and use of supercomputers and whether federal agencies' proposed paths for advancing our supercomputing capabilities are adequate to maintain or regain the U.S. lead.

As we all know, a supercomputer is a broad term for one of the fastest computers currently available. Such computers are typically used for number crunching including scientific simulations, (animated) graphics, analysis of geological data (e.g., in petrochemical prospecting), structural analysis, computational fluid dynamics, physics, chemistry, electronic design, nuclear energy research and meteorology.

Supercomputers are state-of-the-art, extremely powerful computers capable of manipulating massive amounts of data in a relatively short time. They are very expensive and are employed for specialized scientific and engineering applications that must handle very large databases or do a great amount of computation, among them meteorology, animated graphics, fluid dynamic calculations, nuclear energy research and weapon simulation, and petroleum exploration.

Supercomputers are gaining popularity in all corners of corporate America. They are used to analyze vehicle crash test by auto manufacturers, evaluate human diseases and develop treatments by the pharmaceutical industry and test aircraft engines by the aero-space engineers.

It quite evident that supercomputing will become more important to America's commerce in the future. I look forward to working with this committee on its advancement. Again, I wish thank the witnesses for coming here today help us conceptualize this goal.

Chairman BOEHLERT. And just a little history. I can recall back in the early '80's, 1983 to be exact, when I was a freshman and Mr. Hall was an emerging power in the Congress. I sat way down in the front row on the end, and I didn't know what was going on.

But I do remember very vividly the testimony of a Nobel Laureate, Dr. Ken Wilson, who at that time was at Cornell University. And he told us that the typical graduate student in the United Kingdom or Japan or Germany—the typical graduate student had greater access to the latest in computer technology than did he, a young Nobel Laureate, you know, one of the great resources of our nation.

And he argued very forcibly and very persuasively for the Federal Government to get more actively involved. And boy, that—it was like the light bulb going on. I didn't even understand it, and I am not quite sure I do yet, this—the—all of the intricacies of this supercomputer technology, but I do remember then being a champion from day one getting this supercomputer initiative going for America. And in '85, NSF set up the Supercomputing Centers and—at Carnegie Mellon and Cornell and others—and boy, did we just go forward, leapfrog ahead. We did wonderful things.

You know what? A little lethargy is setting in and I am getting concerned. I am deeply concerned, and I mention in my opening statement about five times, I should have mentioned it about 55 times, the importance of a well coordinated federal response to this issue. I don't want to be second to anybody, neither does Mr. Hall, neither do any Members of this committee.

We have an opportunity, and we are going to seize it. And so we are looking at all of you as resources for this committee. You are all very distinguished in your own way. You are very knowledgeable. You will share with us, and hopefully we will get a few more light bulbs turned on up here. And we can go forward together. There is an awful lot at stake.

And so I look forward to your testimony, and I hope that you will sign on here and now. Some of you have no choice. You have to, right? But sign on here and now to work cooperatively with us, because there is so much at stake.

With that, let me introduce our witnesses.

Witness consist—list consists of Dr. Raymond Orbach, Director of the Office of Science, Department of Energy. Dr. Orbach, good to have you back. Dr. Peter A. Freeman, Assistant Director, Computer and Information Science and Engineering Directorate at the National Science Foundation. It is good to see you once again. Dr. Daniel A. Reed, Director, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign. Dr. Reed. And Mr. Vincent Scarafino, Manager, Numerically Intensive Computing, Ford Motor Company. Mr. Scarafino, glad to have you here.

We would ask all of our witnesses to try to summarize your complete statement, because we all have the benefit of the complete statement. And we will read those statements very carefully, but try to summarize in five minutes or so. I am not going to be arbitrary. This is too darn important to restrict your expert input to 300 seconds, but I would ask you to be close to the five minutes. And then we can have a good exchange in the dialogue. And hopefully we will all profit from this little exercise we are engaged in.

Dr. Orbach.

STATEMENT OF DR. RAYMOND L. ORBACH, DIRECTOR, OFFICE OF SCIENCE, DEPARTMENT OF ENERGY

Dr. ORBACH. Chairman Boehlert and Ranking Member Hall, Members of the Committee, I commend you for holding this hearing. And I deeply appreciate the opportunity to testify on behalf of the Office of Science at the Department of Energy on a subject of central importance to this Nation, as you have both outlined, our need for advanced supercomputing capability. Through the efforts of the DOE's Office of Science and other federal agencies, we are working to develop the next generation of advanced scientific computational capacity, a capability that supports economic competitiveness and America's scientific enterprise.

The Bush Administration has forged an integrated and unified interagency road map to the critical problems that you have asked us to address today. In my opening statement, I would like to briefly address the four specific questions that the Committee has asked of me, and more extensive answers are contained in my written testimony.

The first question that the Committee addressed to me concerned the development path for a next-generation advanced scientific computer and whether the Office of Science path differed from that of the National Nuclear Security Agency, the NNSA. The NNSA has stewardship responsibilities and the computer architectures, which they have used, are well suited for those needs. And indeed, they led the way in the massively parallel machine development.

However, those machines operate at only something like five to 10 percent efficiency when applied to many problems of scientific interest and also industrial interest. And that reduces the efficiencies of these high peak speed machines. Other architectures have shown efficiencies closer to 60 percent for some of the physical problems that science and industry must address.

We are working with NNSA as partners to explore these alternatives, which we believe, will be important to both of our areas of responsibility. For example, the Office of Science will be exploring computer architectures that may be of value for magnetic fusion to biology. NNSA is working with us as a partner to explore equations of state under high pressures and extreme temperatures, which, of course, is critical to stewardship issues.

The second question that I was asked was are we cooperating with DARPA, the Defense Advanced Research Project Agency, and how does that relationship work. DARPA has historically invested in new architectures, which have been and are of great interest to the Office of Science. We are—we have a memorandum of understanding [MOU] currently under review between the Defense Department and the Department of Energy that establishes a framework for cooperation between DARPA and the DOE, including both the Office of Science and NNSA.

The MOU will cover high-end computation performance evaluation, development of benchmarks, advanced computer architecture evaluation, development of mathematical libraries, and system software. This will bring together the complementary strengths of each agency. We will be able to draw on DARPA's strengths in advanced computer architectures and they on our decades of experience in

evaluating new architectures and transforming them into tools for scientific discovery.

The third question you asked was to what extent will the Office of Science be guided by the recommendations of the High-End Computing Revitalization Task Force and how will we contribute to the OSTP, Office of Science and Technology Policy, plan to revitalize high-end computation. The formation of the High-End Computation Revitalization Task Force by OSTP emphasizes the importance that the Administration places on the need for a coordinated approach to strengthening high-end computation. A senior official of the Office of Science is co-chairing that task force, and it includes many representatives from across the government and industry.

Many of the task force findings and plans are actually based on Office of Science practices in advanced computing and simulation. We are working very closely with NNSA, the Department of Defense, NASA, the National Science Foundation, and the National Institutes of Health to assess how best to coordinate and leverage our agency's high-end computation investments now and in the future. We expect to play a major role in executing the plans that emerge from this task force in partnership with the other agencies and under the guidance of the President's Science Advisor.

The last question you asked was how are the advanced computational needs of the scientific community and of the private sector diverging and how does that affect advanced computing development programs at the Office of Science. I don't believe there is a major divergence between the needs of the scientific community and those of industry's design engineers. The apparent divergence stems from the dominance of computer development by the needs of specific commercial applications: payroll, management information systems, and web servers.

I will defer to Mr. Scarafino on this, but my own discussions with industry leaders suggest that the type of computer architectures that would meet the needs of the Office of Science would also support their requirements. It would give them the ability to create virtual prototypes of complex systems, allowing engineers to optimize different design parameters without having to build prototypes. This would reduce the time to market. It would decrease the costs, and it would increase economic competitiveness.

These four questions were central, and I appreciate being asked them and given the opportunity to respond. I am gratified that this committee is intent on enabling us to pursue so many important computational opportunities for the sake of scientific discovery, technological innovation, and economic competitiveness.

I thank you for inviting me, and I will be pleased to take questions.

[The prepared statement of Dr. Orbach follows:]

PREPARED STATEMENT OF RAYMOND L. ORBACH

Mr. Chairman and Members of the Committee, I commend you for holding this hearing—and I appreciate the opportunity to testify on behalf of the Department of Energy's (DOE) Office of Science—on a subject of central importance to this nation: our need for advanced supercomputing capability. Through the efforts of DOE's Office of Science and other federal agencies, we are working to develop the next gen-

eration of advanced scientific computational capability, a capability that supports economic competitiveness and America's scientific enterprise.

As will become abundantly clear in my testimony, the Bush Administration has forged an integrated and unified interagency roadmap to the critical problems you have asked us to address today. No one agency can—or should—carry all the weight of ensuring that our scientists have the computational tools they need to do their job, yet duplication of effort must be avoided. The President, and John Marburger, Office of Science and Technology Policy Director, understand this. That is why all of us here are working as a team on this problem.

* * *

Mr. Chairman, for more than half a century, every President and each Congress has recognized the vital role of science in sustaining this nation's leadership in the world. According to some estimates, fully half of the growth in the U.S. economy in the last 50 years stems from federal funding of scientific and technological innovation. American taxpayers have received great value for their investment in the basic research sponsored by the Office of Science and other agencies in our government.

Ever since its inception as part of the Atomic Energy Commission immediately following World War II, the Office of Science has blended cutting edge research and innovative problem solving to keep the U.S. at the forefront of scientific discovery. In fact, since the mid-1940's, the Office of Science has supported the work of more than 40 Nobel Prize winners, testimony to the high quality and importance of the work it underwrites.

Office of Science research investments historically have yielded a wealth of dividends including: significant technological innovations; medical and health advances; new intellectual capital; enhanced economic competitiveness; and improved quality of life for the American people.

Mr. Chairman and Members of this committee, virtually all of the many discoveries, advances, and accomplishments achieved by the Office of Science in the last decade have been underpinned by advanced scientific computing and networking tools developed by the Office of Advanced Scientific Computing Research (ASCR).

The ASCR program mission is to discover, develop, and deploy the computational and networking tools that enable scientific researchers to analyze, model, simulate, and predict complex phenomena important to the Department of Energy—and to the U.S. and the world.

In fact, by fulfilling this mission over the years, the Office of Science has played a leading role in maintaining U.S. leadership in scientific computation worldwide. Consider some of the innovations and contributions made by DOE's Office of Science:

- helped develop the Internet;
- pioneered the transition to massively parallel supercomputing in the civilian sector;
- began the computational analysis of global climate change;
- developed many of the DNA sequencing and computational technologies that have made possible the unraveling of the human genetic code; and
- opened the door for major advances in nanotechnology and protein crystallography.

* * *

Computational modeling and simulation are among the most significant developments in the practice of scientific inquiry in the latter half of the 20th Century. In the past century, scientific research has been extraordinarily successful in identifying the fundamental physical laws that govern our material world. At the same time, the advances promised by these discoveries have not been fully realized, in part because the real-world systems governed by these physical laws are extraordinarily complex. Computers help us to visualize, to test hypotheses, to guide experimental design, and most importantly to determine if there is consistency between theoretical models and experiment. Computer-based simulation provides a means for predicting the behavior of complex systems that can only be described empirically at present. Since the development of digital computers in mid-century, scientific computing has greatly advanced our understanding of the fundamental processes of nature, e.g., fluid flow and turbulence in physics, molecular structure and

reactivity in chemistry, and drug-receptor interactions in biology. Computational simulation has even been used to explain, and sometimes predict, the behavior of such complex natural and engineered systems as weather patterns and aircraft performance.

Within the past two decades, scientific computing has become a contributor to essentially all scientific research programs. It is particularly important to the solution of research problems that are (i) insoluble by traditional theoretical and experimental approaches, e.g., prediction of future climates or the fate of underground contaminants; (ii) hazardous to study in the laboratory, e.g., characterization of the chemistry of radionuclides or other toxic chemicals; or (iii) time-consuming or expensive to solve by traditional means, e.g., development of new materials, determination of the structure of proteins, understanding plasma instabilities, or exploring the limitations of the "Standard Model" of particle physics. In many cases, theoretical and experimental approaches do not provide sufficient information to understand and predict the behavior of the systems being studied. Computational modeling and simulation, which allows a description of the system to be constructed from basic theoretical principles and the available experimental data, are keys to solving such problems.

Advanced scientific computing is indispensable to DOE's missions. It is essential to simulate and predict the behavior of nuclear weapons, accelerate the development of new energy technologies, and the aid in discovery of new scientific knowledge.

As the lead government funding agency for basic research in the physical sciences, the Office of Science has a special responsibility to ensure that its research programs continue to advance the frontiers of science. All of the research programs in DOE's Office of Science—in Basic Energy Sciences, Biological and Environmental Research, Fusion Energy Sciences, and High-Energy and Nuclear Physics—have identified major scientific questions that can only be addressed through advances in scientific computing. This will require significant enhancements to the Office of Science's scientific computing programs. These include both more capable computing platforms and the development of the sophisticated mathematical and software tools required for large scale simulations.

Existing highly parallel computer architectures, while extremely effective for many applications, including solution of some important scientific problems, are only able to operate at 5–10 percent of their theoretical maximum capability on other applications. Therefore, we have initiated a Next Generation Architecture program to evaluate the effectiveness of various different computer architectures in cooperation with the National Nuclear Security Administration (NNSA) and the Defense Advanced Research Project Agency to identify those architectures which are most effective in addressing specific types of simulations.

To address the need for mathematical and software tools, and to develop highly efficient simulation codes for scientific discovery, the Office of Science launched the Scientific Discovery through Advanced Computing (SciDAC) program. We have assembled interdisciplinary teams and collaborations to develop the necessary state-of-the-art mathematical algorithms and software, supported by appropriate hardware and middleware infrastructure to use terascale computers effectively to advance fundamental scientific research essential to the DOE mission.

These activities are central to the future of our mission. Advanced scientific computing will continue to be a key contributor to scientific research as we enter the twenty-first century. Major scientific challenges exist in all Office of Science research programs that can be addressed by advanced scientific supercomputing. Designing materials atom-by-atom, revealing the functions of proteins, understanding and controlling plasma turbulence, designing new particle accelerators, and modeling global climate change, are just a few examples.

* * *

Today, high-end scientific computation has reached a threshold which we were all made keenly aware of when the Japanese Earth Simulator was turned on. The Earth Simulator worked remarkably well on real physical problems at sustained speeds that have never been achieved before. The ability to get over 25 teraFLOPS in geophysical science problems was not only an achievement, but it truly opened a new world.

So the question before us at today's hearing—"Supercomputing: Is the U.S. on the Right Path"—is very timely. There is general recognition of the opportunities that high-end computation provides, and this Administration has a path forward to meet this challenge.

The tools for scientific discovery have changed. Previously, science had been limited to experiment and theory as the two pillars for investigation of the laws of nature. With the advent of what many refer to as "Ultra-Scale" computation," a third pillar—simulation—has been added to the foundation of scientific discovery. Modern computational methods are developing at such a rapid rate that computational simulation is possible on a scale that is comparable in importance with experiment and theory. The remarkable power of these facilities is opening new vistas for science and technology.

Tradition has it that scientific discovery is based on experiment, buttressed by theory. Sometimes the order is reversed, theory leads to concepts that are tested and sometimes confirmed by experiment. But more often, experiment provides evidence that drives theoretical reasoning. Thus, Dr. Samuel Johnson, in his *Preface to Shakespeare,* writes: "Every cold empirick, when his heart is expanded by a successful experiment, swells into a theorist."

Many times, scientific discovery is counter-intuitive, running against conventional wisdom. Probably the most vivid current example is the experiment that demonstrated that the expansion of our Universe is accelerating, rather than in steady state or contracting. We have yet to understand the theoretical origins for this surprise.

During my scientific career, computers have developed from the now "creaky" IBM 701, upon which I did my thesis research, to the so-called massively parallel processors or MPP machines, that fill rooms the size of football fields, and use as much power as a small city.

The astonishing speeds of these machines, especially the Earth Simulator, allow Ultra-Scale computation to inform our approach to science, and I believe social sciences and the humanities. We are now able to contemplate exploration of worlds never before accessible to mankind. Previously, we used computers to solve sets of equations representing physical laws too complicated to solve analytically. Now we can simulate systems to discover physical laws for which there are no known predictive equations. We can model physical or social structures with hundreds of thousands, or maybe even millions, of "actors," interacting with one another in a complex fashion. The speed of our new computational environment allows us to test different inter-actor (or inter-personal) relations to see what macroscopic behaviors can ensue. Simulations can determine the nature of the fundamental "forces" or interactions between "actors."

Computer simulation is now a major force for discovery in its own right.

We have moved beyond using computers to solve very complicated sets of equations to a new regime in which scientific simulation enables us to obtain scientific results and to perform discovery in the same way that experiment and theory have traditionally been used to accomplish those ends. We must think of high-end computation as the third of the three pillars that support scientific discovery, and indeed there are areas where the only approach to a solution is through high-end computation—and that has consequences.

* * *

American industry certainly is fully conversant with the past, present and prospective benefits of high-end computation. The Office of Science has received accolades for our research accomplishments from corporations such as General Electric and General Motors. We have met with the vice presidents for research of these and other member companies of the Industrial Research Institute. We learned, for example, that GE is using simulation very effectively to detect flaws in jet engines. What's more, we were told that, if the engine flaws identified by simulation were to go undetected, the life cycle of those GE machines would be reduced by a factor of two—and that would cause GE a loss of over $100,000,000.00.

The market for high-end computation extends beyond science, into applications, creating a commercial market for ultra-scale computers. The science and technology important to industry can generate opportunities measured in hundreds of million, and perhaps billions of dollars.

Here are just a few examples:

From General Motors:

"General Motors currently saves hundreds of millions of dollars by using its inhouse high performance computing capability of more than 3.5 teraFLOPS in several areas of its new vehicle design and development processes. These include vehicle crash simulation, safety models, vehicle aerodynamics, thermal and combustion

analyses, and new materials research. The savings are realized through reductions in the costs of prototyping and materials used.

However, the growing need to meet higher safety standards, greater fuel efficiency, and lighter but stronger materials, demands a steady yearly growth rate of 30 to 50 percent in computational capabilities but will not be met by existing architectures and technologies.. . .A computing architecture and capability on the order of 100 teraFLOPS for example would have quite an economic impact, on the order of billions of dollars, in the commercial sector in its product design, development, and marketing."

And from General Electric:

"Our ability to model, analyze and validate complex systems is a critical part of the creation of many of our products and design. Today we make extensive use of high-performance computing based technologies to design and develop products ranging from power systems and aircraft engines to medical imaging equipment. Much of what we would like to achieve with these predictive models is out of reach due to limitations in current generation computing capabilities. Increasing the fidelity of these models demands substantial increases in high-performance computing system performance. We have a vital interest in seeing such improvements in the enabling high-performance computing technologies.. . .In order to stay competitive in the global marketplace, it is of vital importance that GE can leverage advances in high-performance computing capability in the design of its product lines. Leadership in high-performance computing technologies and enabling infrastructure is vital to GE if we wish to maintain our technology leadership."

Consider the comparison between simulations and prototyping for GE jet engines.

For evaluation of a design alternative for the purpose of optimization of a compressor for a jet engine design, GE would require 3.1×10^{18} floating point operations, or over a month at a sustained speed of one teraFLOP, which is today's state-of-the-art. To do this for the entire engine would require sustained computing power of 50 teraFLOPS for the same period. This is to be compared with millions of dollars, several years, and designs and re-designs for physical prototyping.

Opportunities abound in other fields such as pharmaceuticals, oil and gas exploration, and aircraft design.

The power of advance scientific computation is just beginning to be realized. One reason that I have emphasized this so much is that some seem to think that advanced scientific computation is the province of the Office of Science and other federal science agencies and therefore is not attractive to the vendors in this field. I believe that's incorrect. I believe instead that our leading researchers are laying out a direction and an understanding of available opportunities. These opportunities spur markets for high-end computation quite comparable to the commercial market which we have seen in the past but requiring the efficiencies and the speeds which high-end computation can provide.

* * *

Discovery through simulation requires sustained speeds starting at 50 to 100 teraFLOPS to examine problems in accelerator science and technology, astrophysics, biology, chemistry and catalysis, climate prediction, combustion, computational fluid dynamics, computational structural and systems biology, environmental molecular science, fusion energy science, geosciences, groundwater protection, high energy physics, materials science and nanoscience, nuclear physics, soot formation and growth, and more (see http://www.ultrasim.info/doe_docs/).

Physicists in Berkeley, California, trying to determine whether our universe will continue to expand or eventually collapse, gather data from dozens of distant supernovae. By analyzing the data and simulating another 10,000 supernovae on supercomputers (at the National Energy Research Scientific Computing Center or NERSC) the scientists conclude that the universe is expanding—and at an accelerating rate.

I just returned from Vienna, where I was privileged to lead the U.S. delegation in negotiations on the future direction for ITER, an international collaboration that hopes to build a burning plasma fusion reactor, which holds out promise for the realization of fusion power. The United States pulled out of ITER in 1998. We're back in it this year. What changed were simulations that showed that the new ITER design will in fact be capable of achieving and sustaining burning plasma. We haven't created a stable burning plasma yet, but the simulations give us confidence that the experiments which we performed at laboratory scales could be realized in larger machines at higher temperatures and densities.

Looking to the future, we are beginning a Fusion Simulation Project to build a computer model that will fully simulate a burning plasma to both predict and interpret ITER performance and, eventually, assist in the design of a commercially feasible fusion power reactor. Our best estimate, however is that success in this effort will require at least 50 teraFLOPS of sustained computing power.

Advances in scientific computation are also vital to the success of the Office of Science's Genomes to Life program.

The Genomes to Life program will develop new knowledge about how micro-organisms grow and function and will marry this to a national infrastructure in computational biology to build a fundamental understanding of living systems. Ultimately this approach will offer scientists insights into how to use or replicate microbiological processes to benefit the Nation.

In particular, the thrust of the Genomes to Life program is aimed directly at Department of Energy concerns: developing new sources of energy; mitigating the long-term effects of climate change through carbon sequestration; cleaning up the environment; and protecting people from adverse effects of exposure to environmental toxins and radiation.

All these benefits—and more—will be possible as long as the Genomes to Life program achieves a basic understanding of thousands of microbes and microbial systems in their native environments over the next 10 to 20 years. To meet this challenge, however, we must address huge gaps not only in knowledge but also in technology, computing, data storage and manipulation, and systems-level integration.

The Office of Science also is a leader in research efforts to capitalize on the promise of nanoscale science.

In an address to the American Association for the Advancement of Science in February 2002, Dr. John Marburger, Director of the Office of Science and Technology Policy, noted, ". . .[W]e are in the early stages of a revolution in science nearly as profound as the one that occurred early in the last century with the birth of quantum mechanics," a revolution spurred in part by "the availability of powerful computing and information technology."

"The atom-by-atom understanding of functional matter," Dr. Marburger continued, "requires not only exquisite instrumentation, but also the capacity to capture, store and manipulate vast amounts of data. The result is an unprecedented ability to design and construct new materials with properties that are not found in nature.. . .[W]e are now beginning to unravel the structures of life, atom-by-atom using sensitive machinery under the capacious purview of powerful computing."

In both nanotechnology and biotechnology, this revolution in science promises a revolution in industry. In order to exploit that promise, however, we will need both new instruments and more powerful computers, and the Office of Science has instituted initiatives to develop both.

We have begun construction at Oak Ridge National Laboratory on the first of five Nanoscale Science Research Centers located to take advantage of the complementary capabilities of other large scientific facilities, such as the Spallation Neutron Source at Oak Ridge, our synchrotron light sources at Argonne, Brookhaven and Lawrence Berkeley, and semiconductor, microelectronics and combustion research facilities at Sandia and Los Alamos. When complete, these five Office of Science nanocenters will provide the Nation with resources unmatched anywhere else in the world.

To determine the level of computing resources that will be required, the Office of Science sponsored a scientific workshop on *Theory and Modeling in Nanoscience,* which found that simulation will be critical to progress, and that new computer resources are required. As a first step to meeting that need, our Next Generation Architecture initiative is evaluating different computer architectures to determine which are most effective for specific scientific applications, including nanoscience simulations.

There are many other examples where high-end computation has changed and will change the nature of the field. My own field is complex systems. I work in a somewhat arcane area called spin glasses, where we can examine the dynamic properties of these very complex systems, which in fact are related to a number of very practical applications. Through scientific simulation, a correlation length was predicted for a completely random material, a concept unknown before. Simulation led to the discovery that there was a definable correlation length in this completely random system. Our experiments confirmed this hypothesis. Again, insights were created that simply were not possible from a set of physical equations that needed solutions, with observable consequences. There are countless opportunities and examples where similar advances could be made.

* * *

As the Chairman and Members of this committee know, the Bush Administration shares Congress' keen interest in high-end computation for both scientific discovery and economic development. A senior official of the Office of Science is co-chairing the interagency High-End Computing Revitalization Task Force, which includes representatives from across the government and the private sector. We are working very closely with the NNSA, the Department of Defense, the National Aeronautics and Space Administration, the National Science Foundation, and the National Institutes of Health to assess how best to coordinate and leverage our agencies' high-end computation investments now and in the future.

DOE is playing a major role in the task force through the Office of Science and the NNSA, and many of the task force's findings and plans are based on Office of Science practices in advanced computing and simulation.

One of the major challenges in this area is one of metrics. How do we know what we are trying to accomplish, and how can we measure how we're getting there? What are the opportunities? What are the barriers? What should we be addressing as we begin to explore this new world?

Our problem in the past has been that, where we have large computational facilities, we have cut them up in little pieces and the large-scale scientific programs that some researchers are interested in have never really had a chance to develop. There's nothing wrong with our process; it is driven by a peer review system. But for some promising research efforts, there simply have not been enough cycles or there wasn't an infrastructure which would allow large-scale simulations to truly develop and produce the kind of discoveries we hope to achieve.

Recognizing this, the Office of Science has announced that ten percent of our National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory—now at ten teraFLOP peak speed—is going to be made available for grand challenge calculations. We are literally going to carve out 4.5 million processor hours and 100 terabytes of disk space for perhaps four or five scientific problems of major importance. We are calling this initiative INCITE—the Innovative and Novel Computational Impact on Theory and Experiment—and we expect to be ready to proceed with it around August 1, 2003. At that time, we will open the competition to all, whether or not they are affiliated with or funded by DOE.

We are launching the INCITE initiative for two reasons. For one, it's the right thing to do: there are opportunities for major accomplishments in this field of science. In addition, there is also a "sociology" that we need to develop.

Given the size and complexity of the machines required for sustained speeds in the 50 to 100 teraFLOPS regime, the sociology of high-end computation will probably have to change. One can think of the usage of ultra-scale computers as akin to that of our current light sources: large machines used by groups of users on a shared basis. Following the leadership of our SciDAC program, interdisciplinary teams and collaborators will develop the necessary state-of-the-art mathematical algorithms and software, supported by appropriate hardware and middleware infrastructure, to use terascale computers effectively to advance fundamental research in science. These teams will associate on the basis of the mathematical infrastructure of problems of mutual interest, working with efficient, balanced computational architectures.

The large amount of data, the high sustained speeds, and the cost will probably lead to concentration of computing power in only a few sites, with networking useful for communication and data processing, but not for core computation at terascale speeds. Peer review of proposals will be used to allocate machine time. Industry will be welcome to participate, as has happened in our light sources. Teams will make use of the facilities as user groups, using significant portions (or all) of the machine, depending on the nature of their computational requirements. Large blocks of time will enable scientific discovery of major magnitude, justifying the large investment ultra-scale computation will require.

We will open our computational facilities to everyone. Ten percent of NERSC's capability will be available to the entire world. Prospective users will not have to have a DOE contract, or grant, or connection. The applications will be peer reviewed, and will be judged solely on their scientific merit. We need to learn how to develop the sociology that can encourage and then support computation of this magnitude; this is a lot of computer time. It may be the case that teams rather than individuals will be involved. It even is possible that one research proposal will be so compelling that the entire ten percent of NERSC will be allocated to that one research question.

The network that may be required to handle that amount of data has to be developed. There is an ES network which we are involved in, and we are studying whether or not it will be able to handle the massive amounts of data that could be produced under this program.

We need to get scientific teams—the people who are involved in algorithms, the computer scientists, and the mathematicians—together to make the most efficient use of these facilities. That's what this opening up at NERSC is meant to do. We want to develop the community of researchers within the United States—and frankly around the world—that can take advantage of these machines and produce the results that will invigorate and revolutionize their fields of study.

But this is just the beginning.

* * *

As we develop the future high-end computational facilities for this nation and world, it is clearly our charge and our responsibility to develop scientific opportunities for everyone. This has been the U.S. tradition. It has certainly been an Office of Science tradition, and we intend to see that this tradition continues, and not just in the physical sciences.

We are now seeing other fields recognizing that opportunities are available to them. In biology, we are aware that protein folding is a very difficult but crucial issue for cellular function. The time scales that biologists work with can scale from a femto-second to seconds-a huge span of time which our current simulation capabilities are unable to accommodate.

High-performance computing provides a new window for researchers to observe the natural world with a fidelity that could only be imagined a few years ago. Research investments in advanced scientific computing will equip researchers with premier computational tools to advance knowledge and to solve the most challenging scientific problems facing the Nation.

With vital support from this committee, the Congress and the Administration, we in the Office of Science will help lead the U.S. further into the new world of supercomputing.

We are truly talking about scientific discovery. We are talking about a third pillar of support. We are talking about opportunities to understand properties of nature that have never before been explored. That's the concept, and it explains the enormous excitement that we feel about this most promising field.

We are very gratified that this committee is so intent on enabling us to pursue so many important opportunities, for the sake of scientific discovery, technological innovation, and economic competitiveness.

Thank you very much.

Biography for Raymond L. Orbach

Dr. Raymond L. Orbach was sworn in as the 14th Director of the Office of Science at the Department of Energy (DOE) on March 14, 2002. As Director of the Office of Science (SC), Dr. Orbach manages an organization that is the third largest federal sponsor of basic research in the United States and is viewed as one of the premier science organizations in the world. The SC fiscal year 2002 budget of $3.3 billion funds programs in high energy and nuclear physics, basic energy sciences, magnetic fusion energy, biological and environmental research, and computational science. SC, formerly the Office of Energy Research, also provides management oversight of the Chicago and Oak Ridge Operations Offices, the Berkeley and Stanford Site Offices, and 10 DOE non-weapons laboratories.

Prior to his appointment, Dr. Orbach served as Chancellor of the University of California (UC), Riverside from April 1992 through March 2002; he now holds the title Chancellor Emeritus. During his tenure as Chancellor, UC–Riverside grew from the smallest to one of the most rapidly growing campuses in the UC system. Enrollment increased from 8,805 to more than 14,400 students with corresponding growth in faculty and new teaching, research, and office facilities.

In addition to his administrative duties at UC–Riverside, Dr. Orbach maintained a strong commitment to teaching. He sustained an active research program; worked with postdoctoral, graduate, and undergraduate students in his laboratory; and taught the freshman physics course each winter quarter. As Distinguished Professor of Physics, Dr. Orbach set the highest standards for academic excellence. From his arrival, UC–Riverside scholars led the Nation for seven consecutive years in the number of fellows elected to the prestigious American Association for the Advancement of Science (AAAS).

Dr. Orbach began his academic career as a postdoctoral fellow at Oxford University in 1960 and became an assistant professor of applied physics at Harvard University in 1961. He joined the faculty of the University of California, Los Angeles (UCLA) two years later as an associate professor, and became a full professor in 1966. From 1982 to 1992, he served as the Provost of the College of Letters and Science at UCLA.

Dr. Orbach's research in theoretical and experimental physics has resulted in the publication of more than 240 scientific articles. He has received numerous honors as a scholar including two Alfred P. Sloan Foundation Fellowships, a National Science Foundation Senior Postdoctoral Fellowship, a John Simon Guggenheim Memorial Foundation Fellowship, the Joliot Curie Professorship at the Ecole Superieure de Physique et Chimie Industrielle de la Ville de Paris, the Lorentz Professorship at the University of Leiden in the Netherlands, and the 1991–1992 Andrew Lawson Memorial Lecturer at UC–Riverside. He is a fellow of the American Physical Society and the AAAS.

Dr. Orbach has also held numerous visiting professorships at universities around the world. These include the Catholic University of Leuven in Belgium, Tel Aviv University, and the Imperial College of Science and Technology in London. He also serves as a member of 20 scientific, professional, or civic boards.

Dr. Orbach received his Bachelor of Science degree in Physics from the California Institute of Technology in 1956. He received his Ph.D. degree in Physics from the University of California, Berkeley, in 1960 and was elected to Phi Beta Kappa.

Dr. Orbach was born in Los Angeles, California. He is married to Eva S. Orbach. They have three children and seven grandchildren.

Chairman BOEHLERT. Thank you very much, Dr. Orbach.

You noticed the red light was on, and the Chair was generous with the time. As I said, I am not going to be arbitrary. I wish the appropriators were as generous with the funding for your office as we are with time for your views, but we are working continually together on that one.

Dr. Freeman.

STATEMENT OF DR. PETER A. FREEMAN, ASSISTANT DIRECTOR, COMPUTER AND INFORMATION SCIENCE AND ENGINEERING DIRECTORATE, NATIONAL SCIENCE FOUNDATION

Dr. FREEMAN. Good morning, Mr. Chairman, Mr. Hall, and distinguished Members of the Committee. NSF deeply appreciates this committee's long time support and recognizes your special interest in computing, so I am delighted to be here today to discuss those topics with you.

Supercomputing is a field that NSF, as you noted in your opening remarks, has championed for many years. And it is one in which we intend to continue to lead the way. At the same time, we are committed to realizing the compelling vision described in the report of the recent NSF Advisory Panel on Cyberinfrastructure, commonly known as the Atkins Committee. They have forcefully told us, and I quote, "A new age has dawned in scientific and engineering research, " that will be enabled by cyberinfrastructure.

The term "cyberinfrastructure" sounds exotic and is sometimes confused as being something new, but in reality, it is intended to signify a set of integrated facilities and services essential to the conduct of leading edge science and engineering. Cyberinfrastructure must include a range of supercomputers as well as massive storage, high performance networks, databases, lots of software, and above all, highly trained people. An advanced cyberinfrastructure, with supercomputing as an important element, promises to revolutionize research in the 21st Century. The opportunities that are presented to us in this area must be exploited for

the benefit of all of our citizens for their continuing health, security, education, and wealth.

We are committed to a key recommendation of the Atkins Report, namely that NSF, in partnership with other agencies and other organizations, must make significant investments in the creation, deployment, and application of advanced cyberinfrastructure to empower continued U.S. leadership in science and engineering.

To bring about this scientific revolution, we must maintain a broad discourse about cyberinfrastructure. Supercomputers are essential, but without software, networks, massive storage, databases, and trained people all integrated securely, they will not deliver their potential. Further, there are now many areas of science that need this integrated and balanced support or balanced approach more than they need any single element.

My written testimony provides detailed answers to the questions you posed for this hearing, but allow me to summarize them. NSF will most definitely continue its commitment to supercomputing and will provide the support necessary to utilize it in all of science and engineering. Supercomputing capability is an essential component in the cyberinfrastructure vision, and we are committed to expanding this capability in the future.

NSF's recent investments in grid computing underscore the importance of integrating supercomputing within a cyberinfrastructure. Indeed, the first increment of our funding and our terascale efforts was for a supercomputer, which at the time it was installed, was the most powerful open access machine in the research world. This machine will be one of the main resources on the grid that is currently under construction.

The Atkins Report recommended "a 2-year extension of the current PACI [Partnerships for Advanced Computational Infrastructure] cooperative agreements," and we have done that. The panel also recommended "new separately peer-reviewed enabling and application infrastructure" programs. We have identified funding for these efforts and are in the process of working out the details for these new programs. Finally, the Atkins Report recommends that "until the end of the original 10-year lifetime of the PACI program," the centers "should continue to be assured of stable, protected funding to provide the highest-end computing resources." NSF is currently gathering community input on how best to structure future management of these resources, and will be announcing our specific plans in the coming months.

With regard to coordination, NSF has been an active participant in all four of the subgroups of the OSTP planning activity for high-end computing. We are coordinating closely at all levels with OSTP and our sister agencies to make sure that the recommendations of the task force are carried forward in an effective manner.

NSF does not see a divergence in the needs of industry and those of the research community. As researchers push the boundaries in their work, the results are often quickly picked up by industry. In other areas, industry has had to tackle problems first, such as data mining, and now those techniques are becoming useful in scientific research. This symbiotic relationship has worked very well in the past, and we intend to continue it in the future.

Mr. Chairman, let me close by providing an answer to the overall question of whether we are on the right path in supercomputing. My answer is yes; if we keep in perspective that supercomputers must be embedded in a cyberinfrastructure that also includes massive storage, high-performance networks, databases, lots of software, well-trained people, and that the entire ensemble be open to all scientists and engineers.

Thank you, and I will be glad to answer any questions you may have.

[The prepared statement of Dr. Freeman follows:]

PREPARED STATEMENT OF PETER A. FREEMAN

Good morning, Mr. Chairman and Members of the Committee. I am Dr. Peter Freeman, Assistant Director of the NSF for CISE.

Introduction

I am delighted to have the opportunity to testify before you this morning and to discuss the topic *Supercomputing: Is the U.S. on the Right Path?* Supercomputers are an extremely important national resource in many sectors of our society, and for decades, these resources have yielded astounding scientific breakthroughs. Supercomputing is a field that NSF has championed and supported for many years and it is one in which we will continue to lead the way.

There seems to be some confusion in the scientific community as to NSF's commitment to High-End Computing (HEC) which is the current term being used for "supercomputing." I want to clear the air this morning. Before I briefly summarize my written testimony, *let me state unequivocally that NSF remains absolutely committed to providing researchers the most advanced computing equipment available and to sponsoring research that will help create future generations of computational infrastructure, including supercomputers.*

At the same time, we are committed to realizing the compelling vision described in the report of the NSF Advisory Panel on Cyberinfrastructure, commonly known as the Atkins Committee—that "a new age has dawned in scientific and engineering research, pushed by continuing progress in computing, information and communications technology." This cyberinfrastructure includes, and I quote, "not only high-performance computational services, but also integrated services for knowledge management, observation and measurement, visualization and collaboration."

The scientific opportunities that lie before us in many fields can *only* be realized with such a cyberinfrastructure. Just as supercomputing promised to revolutionize the conduct of science and engineering research several decades ago, and we are seeing the results of that promise today, so does an advanced cyberinfrastructure promise to revolutionize the conduct of science and engineering research and education in the 21st century. The opportunities that a balanced, state-of-the-art cyberinfrastructure promises *must* be exploited for the benefit of all of our citizens—for their continuing health, security, education, and wealth.

To be clear, we are committed to what the Atkins Report and many others in the community, both formally and informally, are saying: That NSF, in partnership with other public and private organizations, must make investments in the creation, deployment and application of cyberinfrastructure in ways that radically empower all science and engineering research and allied education. . .thereby empowering what the Atkins Report defines as "a revolution."

Cyberinfrastructure—with HEC as an essential component—can bring about this true revolution in science and engineering. It promises great advances for all of the areas of our society served by science and engineering, but it will be realized only if we stay focused on the value of all components of cyberinfrastructure.

Supercomputers have been one of the main drivers in this revolution up to now because of the continuing evolution of computing technology. Computers were initially developed to deal with pressing, numerical computations and they will continue to be extremely important. In recent years, however, many scientific advances have been enabled not only by computers, but by the great expansion in the capacity of computer storage devices and communication networks, coupled now with rapidly improving sensors.

There are now many examples of revolutionary scientific advances that can only be brought about by utilizing other components of cyberinfrastructure in combination with HEC. This necessary convergence means that we must maintain a broad discourse about cyberinfrastructure. As we set about building and deploying an ad-

vanced cyberinfrastructure, we will ensure that the HEC portion remains an extremely important component in it.

History of NSF Support for Supercomputing

NSF has been in the business of supporting high performance computation in the form of centers since the establishment of the first Academic Computing Centers in the 1960's. As computers became increasingly powerful, they were later designated to be "supercomputers." In the mid-1980's, NSF created the first supercomputing centers for the open science community. A decade later, support was established through the Partnerships for Advanced Computational Infrastructure (PACI) program. In 2000, a parallel activity that is now known as the Extensible Terascale Facility was initiated. (There is much more to this history that is available if desired.) Beginning in FY 2005, support will be provided through cyberinfrastructure program(s) currently under development.

Over time, technological innovations have led to movement away from the use of the term "supercomputing centers" since it inadequately describes the full measure and promise of what is being done at such centers, and what is at stake. The idea of the "*super*computer" lives on as a legacy, however a more accurate title for this kind of infrastructure would be High-performance Information Technology (HIT) Centers.

NSF currently supports three major HIT Centers: the San Diego Supercomputer Center (SDSC), the National Center for Supercomputing Applications (NCSA), and the Pittsburgh Supercomputer Center (PSC). These centers have been participating in this enterprise since the days of the first supercomputer centers in the early 1980's. They have evolved steadily over the past quarter century and now represent special centers of talent and capability for the science community broadly and the Nation at large.

In the last six years, NSF's support for these HIT Centers has been provided predominantly through the PACI program. More recently, a new consortium of HIT Centers has emerged around the new grid-enabled concept of the Extensible Terascale Facility (ETF).

Current Activities

In order to describe NSF's current activities in the area of supercomputing, I'd like to respond directly to the following questions formulated by Chairman Boehlert and the Committee on Science. (Note: Italicized and numbered statements are drawn verbatim from Chairman Boehlert's letter of invitation.)

1. *Some researchers within the computer science community have suggested that the NSF may be reducing its commitment to the supercomputer centers. Is this the case?*

NSF is most definitely **not** reducing its commitment to supercomputing.

For several decades the agency has invested millions of taxpayer dollars in the development and deployment of a high-end computational infrastructure. These resources are made widely available to the science and engineering research and education community. The agency is not reducing its commitment to such efforts. In fact, leading-edge supercomputing capabilities are an essential component in the cyberinfrastructure and, in line with the recommendations of the Advisory Panel on Cyberinfrastructure, we are committed to expanding such capabilities in the future.

1. *(cont'd.) To what extent does the focus on grid computing represent a move away from providing researchers with access to the most advanced computing equipment?*

The term "grid computing" is ambiguous and often misused. It sometimes is used to signify a single computational facility composed of widely separated elements (nodes) that are interconnected by high-performance networks and that are operating in a manner that the user sees a single "computer." This is a special case of the more general concept of a set of widely separated computational resources of different types, which can be accessed and utilized as their particular capabilities are needed.

While still experimental at this stage, grid computing promises to become the dominant modality of High-performance IT (and, eventually, of commodity computing). One need only think about the World Wide Web (WWW) to understand the compelling importance of grid computing. In the WWW one is able from a single terminal (typically a desk-top PC) to access many different databases, on-line services, even computational engines today. Imagine now that the nodes that are accessible in this way are HECs with all of their computational power; or massive data stores that can be manipulated and analyzed; or sophisticated scientific instruments to ac-

quire data; or any of a dozen other foreseen and unforeseen tools. With this vision, perhaps one can understand the promise of grid computing.

NSF's recent investments in grid computing, through the ETF, should not be seen as a reduction in the agency's commitment to HEC. Rather, it underscores the importance of HEC integrated into a broad cyberinfrastructure. Indeed, the first increment of ETF funding was for a HEC machine, which at the time it was installed at the PSC, was the most powerful open access machine in the research world (and it is, three years later, still number 9 on the Top 500 list). This machine is one of the main resources on the ETF grid. While NSF may not always have the world's fastest machine, we will continue to provide a range of supercomputing systems that serve the ever increasing and changing needs of science and engineering.

At the same time, the ETF investment in grid computing is not the ONLY investment the agency has recently made in HEC. In fact, HEC upgrades at NCSA and SDSC during FY 2003 and 2004 are expected to fund the acquisition of an additional 20 Teraflops of HEC capability.

NSF's unwavering commitment is to continuously advance the frontier, and grid computing is widely acknowledged to represent the next frontier in computing. In short, the ETF represents our commitment to innovation at the computing frontier.

2. *What are the National Science Foundation's (NSF's) plans for funding the supercomputer centers beyond fiscal year 2004? To what extent will you be guided by the recommendation of the NSF Advisory Panel on Cyberinfrastructure to maintain the Partnerships for Advanced Computational Infrastructure, which currently support the supercomputer centers?*

NSF's plans for funding supercomputer centers beyond FY 2004 are very much guided by the recommendations of the NSF Advisory Panel on Cyberinfrastructure as described below.

- In its report, the Panel recommended "a two-year extension of the current PACI co-operative agreements." The National Science Board approved the second one-year extension of the PACI cooperative agreements at the May 2003 meeting.
- The Panel also recommended that ". . .the new separately peer-reviewed enabling and application infrastructure would begin in 2004 or 2005, after the two-year extensions of the current cooperative agreements."

  The President requested $20 million for NSF in FY 2004 for activities that will focus on the development of cyberinfrastructure, including enabling infrastructure (also known as enabling technology). This increased investment in enabling technology will strengthen the agency's portfolio of existing awards, and as the Panel recommended, awards will be identified through merit-review competition.

  In addition, NSF will also increase its investments in applications infrastructure (also known as applications technology) by drawing upon interdisciplinary funds available in the FY 2004 ITR priority area activity. Again, the most promising proposals will be identified using NSF's rigorous merit review process.
- Finally, the Panel's Report recommends that "After these two years, until the end of the original 10-year lifetime of the PACI program, the panel believes that" NCSA, SDSC and PSC "should continue to be assured of stable, protected funding to provide the highest-end computing resources."

  Accordingly, and in planning how to provide such support while positioning the community to realize the promise of cyberinfrastructure, NSF has held a series of workshops and town hall meetings over the course of the past two months to gather community input. Informed by community input, support for SDSC, NCSA and PSC will be provided through new awards to be made effective the beginning of FY 2005.

NSF has also committed to providing support for the management and operations of the Extensible Terascale Facility through FY 2009; this includes support for SDSC, NCSA and PSC who are partners in the ETF.

3. *To what extent will NSF be guided by the recommendations of the High-End Computing Revitalization Task Force? How will NSF contribute to the Office of Science and Technology Policy plan to revitalize high-end computing?*

NSF has been an active participant in all four of the subgroups of the OSTP current planning activity called HEC–RTF. The final report is not finished, but we continue to coordinate closely at all levels with OSTP and its sister agencies to make sure that the recommendations of the Task Force are carried forward in a produc-

tive and effective manner. NSF's traditional role as a leader in innovating new HEC computational mechanisms and applications, and in ensuring that there are appropriate educational programs in place to train scientists and engineers to use them, must be integral to any efforts to increase HEC capabilities for the Nation.

We are now planning our request for the FY 2005 budget, and cyberinfrastructure, including HEC, is likely to be a major component in it. We intend to continue to have more than one high-end machine for the NSF community to access and to invest in needed HEC capabilities as noted above.

4. To what extent are the advanced computational needs of the scientific community and of the private sector diverging? What is the impact of any such divergence on the advanced computing programs at NSF?

I don't believe that the advanced computational "needs" of the science community and the private sector are diverging. In fact, I believe that the growing scientific use of massive amounts of data parallels what some sectors of industry already know. In terms of HEC, it is clear that both communities need significantly faster machines to address those problems that can only be solved by massive computations.

For several decades, NSF has encouraged its academic partners in supercomputing, including NCSA, SDSC and PSC, to develop strong relationships with industry. The initial emphasis was on "supercomputing." And many of the industrial partners at the centers learned about supercomputing in this way, and then started their own internal supercomputer centers.

When Mosaic (the precursor to Netscape) was developed at NCSA, industry was able to rapidly learn about and exploit this new, revolutionary technology. As noted above, grid computing, which is being innovated at NCSA, SDSC and PSC and their partners today, is already being picked up by industry as a promising approach to tomorrow's computing problems.

As researchers push the boundaries in their work, the results (and means of obtaining those results) are often quickly picked up by industry. Conversely, in some areas industry has had to tackle problems (such as data mining) first and now those techniques are becoming useful in scientific research. We intend to continue the close collaboration that has existed for many years.

Conclusion

Mr. Chairman, I hope that this testimony dispels any doubt about NSF's commitment to HEC and the HIT Centers that today provide significant value to the science and engineering community.

I hope that I have also been able to articulate that the cyberinfrastructure vision eloquently described in the Atkins report includes HEC and other advanced IT components. This cyberinfrastructure will enable a true revolution in science and engineering research and education that can bring unimagined benefits to our society.

NSF recognizes the importance of HEC to the advanced scientific computing infrastructure for the advancement of science and knowledge. We are committed to continuing investments in HEC and to developing new resources that will ensure that the United States maintains the best advanced computing facilities in the world. We look forward to working with you to ensure that these goals are fulfilled.

Thank you for the opportunity to appear before you this morning.

Biography for Peter A. Freeman

Peter A. Freeman became Assistant Director for the Computer and Information Science and Engineering Directorate (CISE) on May 6, 2002.

Dr. Freeman was previously at Georgia Institute of Technology as professor and founding Dean of the College of Computing since 1990. He served in that capacity as the John P. Imlay, Jr. Dean of Computing, holding the first endowed Dean's Chair at Georgia Tech. He also served as CIO for the campus for three years. In addition, as a general officer of the campus, he was heavily involved in planning and implementing a wide range of activities for the campus including a successful $700M capital campaign and the Yamacraw Economic Development Mission. He was in charge of the FutureNet Project, part of the campus technology preparations for the 1996 Olympic Village, that resulted in a very high-performance and broad campus network. In 1998, he chaired the Sam Nunn NationsBank Policy Forum on Information Security which lead to the creation of the Georgia Tech Information Security Center, one of the first comprehensive centers in the country focused on information security.

During 1989–90 Dr. Freeman was Visiting Distinguished Professor of Information Technology at George Mason University in Fairfax, Virginia, and from 1987 to 1989

he served as Division Director for Computer and Computation Research at the National Science Foundation. He served on the faculty of the Department of Information and Computer Science at the University of California, Irvine, for almost twenty years before coming to Georgia Tech.

He co-authored *The Supply of Information Technology Workers in the United States* (CRA, 1999) and authored *Software Perspectives: The System is the Message* (Addison Wesley, 1987), *Software Systems Principles* (SRA, 1975), and numerous technical papers. In addition, he edited or co-edited four books including, *Software Reusability* (IEEE Computer Society, 1987), and *Software Design Techniques, 4th edition* (IEEE Press, 1983). He was the founding editor of the *McGraw-Hill Series in Software Engineering and Technology,* has served on several editorial boards and numerous program committees, and was an active consultant to industry, academia, and government.

Dr. Freeman was a member of the Board of Directors of the Computing Research Association (1988–2002), serving as Vice-Chair and Chair of the Government Affairs Committee. He was a member of select review committees of the IRS and FAA Airtraffic Control modernization efforts, and has served on a variety of national and regional committees. While at NSF, he helped formulate the High-Performance Computing and Communications Initiative of the Federal Government.

Dr. Freeman is a Fellow of the IEEE (Institute for Electrical and Electronics Engineers), AAAS (American Association for the Advancement of Science), and the ACM (Association for Computing Machinery). He received his Ph.D. in computer science from Carnegie-Mellon University in 1970, his M.A. in mathematics and psychology from the University of Texas at Austin in 1965, and his B.S. in physics from Rice University in 1963. His research and technical expertise has focused on software systems and their creation. His earliest work (1961–63) involved developing advanced scientific applications in the days before there were operating systems and other support software. This led him to design and build one of the earliest interactive time-sharing operating systems (1964) and ultimately to early work applying artificial intelligence to the design process for software (1965–75). This culminated with the publication of his first book, *Software System Principles* (SRA, 1975).

After a short stint teaching overseas for the United Nations, he focused his work on software engineering, ultimately being recognized for this early work by being elected a Fellow of the IEEE. Along with Prof. A.I. Wasserman, he developed one of the first software design courses (taken by thousands of industry practitioners) and published a highly popular text that served as a first introduction to software engineering. His research during this period focused on reusable software, especially using formal transformation systems. That work has resulted in several startup companies.

Since 1987 when he was "loaned" by the University of California to the National Science Foundation, he has focused his attention on national policy and local action intended to advance the field of computing. In addition to his many activities as Dean at Georgia Tech, he headed an NSF-funded national study of the IT worker shortage (http://www.cra.org/reports/wits/cra.wits.html), started an active group for Deans of IT& Computing, and published several papers relative to future directions of the field.

Chairman BOEHLERT. Thank you very much, Dr. Freeman.

The Chair recognizes Mr. Johnson of Illinois.

Mr. JOHNSON. Thank you, Mr. Chairman. I appreciate the opportunity to introduce our next guest, and I certainly commend the Chairman and the Committee on the amble of the hearing as well as the quality of the witnesses that we have here.

It gives me a tremendous amount of pleasure to represent—or to introduce to you Dr. Daniel A. Reed, who is the Director of the National Center for Computing—Supercomputing Applications at the University of Illinois, my home university, the University of Illinois at Urbana-Champaign. Dr. Reed, who is obviously here with us today, serves as the Director of the National Computational Science Alliance, the National Center for Computing—Supercomputing Applications at the University of Illinois at Urbana-Champaign. Dr. Reed is one of two principal investigators and the chief architect for the NSF TeraGrid project to create a U.S. national infrastructure for grid computing. He is an incoming member of the Presi-

dent's Information Technology Advisory Committee and was formerly, from 1996 to 2001, I believe, Dr. Reed was the head of the University of Illinois Computer Science Department and a co-leader of the National Computational Science Alliance Enabling Technology team for three years.

He brings a tremendous amount of pride to our tremendous university, which I think is recognized worldwide as one of the leaders in this area. And the impact that Dr. Reed and our center has had on the computer science field, not only here but around the country and the world, for that matter, is really indescribable.

So with those brief introductory remarks and my appreciation to the Chair and the Committee as well as Dr. Reed, it gives me a great deal of pleasure to introduce to the Committee and—for his testimony, Dr. Daniel A. Reed.

Dr. REED. Thank you, Congressman Johnson.

Chairman BOEHLERT. Dr. Reed, with a glowing introduction.

STATEMENT OF DR. DANIEL A. REED, DIRECTOR, NATIONAL CENTER FOR SUPERCOMPUTING APPLICATIONS, UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

Dr. REED. Good morning, Mr. Chairman, and Members of the Committee. In response to your questions, I would like to make three points today regarding the status and future of advanced scientific computing.

First, the National Science Foundation has played a critical role in providing the high-end computational infrastructure for the science and engineering community, a role that really must continue. Via this program, the NSF Supercomputing Centers have provided community access to high-end computers that, as you noted earlier, had previously been available only in restricted cases and to researchers at National Laboratories. This access is by national peer review, which ensures that the most promising proposals from across science and engineering benefit from high-end computing support.

NSF's planned cyberinfrastructure will play a critical role in accelerating scientific discovery, as Dr. Freeman noted, by coupling scientific data archives, instruments, and computing resources via distributed grids. The important issue, though, and the one we are discussing today, is a relative level of investment in distributed cyberinfrastructure and high-end computing.

Grids, which were pioneered by the high-end computing community, are not a substitute for high-end computing. Many problems of national importance can only be solved by tightly coupled high-end computing systems.

This brings me to my second point: the challenges before us. On behalf of the Interagency High-End Computing Revitalization Task Force, I recently chaired a workshop on high-end computing opportunities and needs. And at the workshop, researchers made compelling cases for sustained performance levels of 25 to 100 teraflops to enable new scientific discoveries. By comparison, the aggregate peak performance of the system's now available in NSF's three Supercomputing Centers and the Department of Energy's National Energy Research Scientific Computing Center, or NERSC, is roughly 25 teraflops. Simply put, there are neither enough high-end com-

puting systems available nor are there capabilities adequate to address critical research opportunities. And planned procurements, the ones on the horizon, will not fully address that shortfall.

At the workshop, researchers also cited the difficulty in achieving high performance on these machines. In the 1990's, as it was mentioned earlier, the U.S. High-Performance Computing Communications program supported the development of several new machines. But in retrospect, we did not learn a critical lesson, a lesson from the 1970's with vector systems, namely the need for long-term, sustained investment in both hardware and software. In the 1990's, we under-invested in software, and we expected innovative research projects to yield robust, mature software for use by the science community in only two to three years.

In addition to that, we are now extremely dependent on hardware derived, in my opinion, too narrowly from the commercial market. Many, though I emphasize not all, scientific applications and, indeed, several critical to weapons design, signals intelligence, and cryptanalysis are not well matched to the capabilities of commercial systems. I believe new high-end computing designs will be needed to address some of these scientific applications, both fully custom high-end designs, as well as more appropriate designs based on commodity components.

This leads me to my third and final point: appropriate levels of cooperation and support for high-end computing. I believe we must change the model of development and deployment for high-end systems in the U.S. if we are to sustain the leadership needed for continued scientific discovery and national security. The Japanese Earth System Simulator is a wake-up call as it highlights the critical importance of industry/government collaboration and sustained investment.

To sustain U.S. leadership, I believe we must pursue two concurrent paths. First, we must continue to acquire and deploy high-end systems, but at larger scale if we are to satisfy the unmet demands of the open research community. As many have noted, many community reports and panels have made this point, including the recent NSF cyberinfrastructure report. NSF should build on its high-end computing successes and implement the high-end computing recommendations of the cyberinfrastructure report, increasing and sustaining the long-term investment in high-end computing hardware, software, and the support staff needed to catalyze scientific discovery.

But the need for high-end computing is both broad and deep. It has both research components as well as mission applications. Hence, high-end computing system deployment should not be viewed as an interagency competition, but rather as an unmet need that requires aggressive responses from many agencies. Their complimentary roles in each agency has a place to play there.

Second, and concurrently, we must begin a coordinated R&D effort to create systems that are better matched to scientific application needs. I believe we must fund the construction of large-scale prototypes with balanced exploration of new hardware and software models driven critically by scientific application requirements. This cycle of coordinated prototyping assessment and commer-

cialization must be a long-term and sustaining investment. It cannot be a one-time crash program.

Thank you very much for your time and attention. I will, at the appropriate time, be happy to answer questions.

[The prepared statement of Dr. Reed follows:]

PREPARED STATEMENT OF DANIEL A. REED

Good morning, Mr. Chairman and Members of the Committee. Thank you very much for granting me this opportunity to comment on appropriate paths for scientific computing. I am Daniel Reed, Director of the National Center for Supercomputing Applications (NCSA), one of three NSF-funded high-end computing centers. I am also a researcher in high-performance computing and a former head of the Department of Computer Science at the University of Illinois.

In response to your questions, I would like to make three points today regarding the status and future of advanced scientific computing.

1. Success and Scientific Opportunities

First, the National Science Foundation has played a critical role in providing high-end computational infrastructure for the broad university-based science and engineering community, a critical role that must continue. NCSA began with NSF funding for the supercomputer centers program in the 1980s. Via this program and its successors, the NSF supercomputing centers have provided community access to high-end computers that previously had been available only to researchers at national laboratories. These supercomputing investments have not only enabled scientific discovery, they have also catalyzed development of new technologies and economic growth.

The Internet sprang from early DARPA investments and from funding for NSFNet, which first connected NSF's supercomputing centers and provided open access to high-end computing facilities. NCSA Mosaic, the first graphical web browser, which spawned the web revolution, grew from development of tools to support collaboration among distributed scientific groups. Research by NCSA and the other supercomputing centers was instrumental in the birth of scientific visualization—the use of graphical imagery to provide insight into complex scientific phenomena. Via ONR-funded security center, NCSA is creating new cybersecurity technologies to safeguard the information infrastructure of our nation and our military forces.

Today, the NSF supercomputing centers and their Partnerships for Advanced Computational Infrastructure (PACI) program and Extended Terascale Facility (ETF) partners are developing new tools for analyzing and processing the prodigious volumes of experimental data being produced by new scientific instruments. They are also developing new Grid technologies that couple distributed instruments, data archives and high-end computing facilities, allowing national research collaborations on an unprecedented scale.

Access to high-end computing facilities at NCSA, the San Diego Supercomputing Center (SDSC) and the Pittsburgh Supercomputing Center (PSC) is provided by national peer review, with computational science researchers reviewing the proposals for supercomputing time and support. This process awards access to researchers without regard to the source of their research funds. The three centers support researchers with DOE, NIH and NASA awards, among others. This ensures that the most promising proposals from across the entire science and engineering community gain access to high-end computing hardware, support staff and software. Examples include:

- Simulation of cosmological evolution to test basic theories of the large-scale structure of the universe
- Quantum chromodynamics simulations to test the Standard Model of particle physics
- Numerical simulation of weather and severe storms for accurate prediction of severe storm tracks
- Climate modeling to understand the effects of climate change and assess global warming
- Studying the dynamics and energetics of complex macromolecules for drug design
- Nanomaterials design and assessment
- Fluid flow studies to design more fuel-efficient aircraft engines

From these, let me highlight just two examples of scientific exploration enabled by high-end computing, and the limitations on computing power we currently face.

A revolution is underway in both astronomy and high-energy physics, powered in no small part by high-end scientific computing. New data, taken from high-resolution sky surveys, have exposed profound questions about the nature and structure of the universe. We now know that the overwhelming majority of the matter in the universe is of an unknown type, dubbed "dark matter," that is not predicted by the Standard Model of physics. Observations also suggest that an unknown "dark energy" is causing the universe to expand at an ever-increasing rate. Both of these discoveries are profound and unexpected.

Large-scale simulations at NCSA and other high-end computing centers are being used to investigate models of the evolution of the universe, providing insight in the viability of competing theories. The goal of these simulations is to create a computational "universe in a box" that evolves from first principles to conditions similar to that seen today. *These cosmological studies also highlight one of the unique capabilities of large-scale scientific simulation—the ability to model phenomena where experiments are otherwise not possible.*

As a second example, the intersecting frontiers of biology and high-end computing are illuminating biological processes and medical treatments for disease. We have had many successes and others are near. For example, given large numbers of small DNA fragments, high-end computing enabled "shotgun sequencing" approaches for assembling the first maps of the human genome. High-end computing has also enabled us to understand how water moves through cell walls, how blood flow in the heart can disrupt the plaque that causes embolisms, and how new drugs behave.

Unraveling the DNA code for humans and organisms has enabled biologists and biomedical researchers to ask new questions, such as how genes control protein formation and cell regulatory pathways and how different genes increase the susceptibility to disease. Biophysics and molecular dynamics simulations are now being used to study the structure and behavior of biological membranes, drug receptors and proteins. In particular, understanding how proteins form three-dimensional structures is central to designing better drugs and combating deadly diseases such as HIV and SARS. *Today's most powerful high-end computing systems can only simulate microseconds of the protein folding process; complete folding takes milliseconds or more. Such large-scale biological simulations will require vast increases in computing capability, perhaps as much as 1000 times today's capability.*

Simply put, we are on the threshold of a new era of scientific discovery, enabled by computational models of complex phenomena. From astronomy to zoology, high-end computing has become a peer to theory and experiment for exploring the frontiers of science and engineering.

This brings me to my second point: the challenges before us.

2. Challenges

Although large-scale computational simulation has assumed a role equal to experiment and theory in the scientific community, within that community, we face critical challenges. There is a large and unmet demand for access to high-end computing in support of basic scientific and engineering research. There are neither enough high-end computing systems available nor are their capabilities adequate to address fully the research challenges and opportunities. This view is supported by recent workshops, reports and surveys, including the NSF report on high-end computing and cyberinfrastructure, the DOD integrated high-end computing report, and the DOE study on a science case for large-scale simulation.

On behalf of the interagency High-End Computing Revitalization Task Force (HECRTF), I recently chaired a workshop to gain community input on high-end computing opportunities and needs. At the workshop, researchers from multiple disciplines made compelling cases for sustained computing performance of 50–100X beyond that currently available. Moreover, researchers in every discipline at the HECRTF workshop cited the difficulty in achieving high, sustained performance (relative to peak) on complex applications to reach new, important scientific thresholds. Let me cite just a few examples to illustrate both the need and our current shortfall.

In high-energy physics, lattice quantum chromodynamics (QCD) calculations, which compute the masses of fundamental particles from first principles, require a *sustained performance* of 20–50 teraflops/second (one teraflop is 10^{12} arithmetic operations per second). This would enable predicative calculations for ongoing and planned experiments. In magnetic fusion research, sustained performance of 20 teraflops/second would allow full-scale tokamak simulations, providing insights into the design and behavior of proposed fusion reactor experiments such as the international ITER project. HECRTF workshop participants also estimated that a sustained performance of 50 teraflops/second would be needed to develop realistic mod-

els of complex mineral surfaces for environmental remediation, and to develop new catalysts that are more energy efficient.

Note that each of these requirements is for sustained performance, rather than peak hardware performance. This is notable for two reasons. First, the aggregate peak performance of the high-end computing systems now available at NSF's three supercomputing centers (NCSA, SDSC and PSC) and DOE's National Energy Research Scientific Computing Center (NERSC) is roughly 25 teraflops.[1] Second, researchers in every discipline at the HECRTF workshop cited the difficulty in achieving high, sustained performance (relative to peak) on complex applications.

Simply put, the Nation's aggregate, open capability in high-end computing is at best equal to the scientific community's estimate of that needed for a single, breakthrough scientific application study. This is an optimistic estimate, as it assumes one could couple all these systems and achieve 100 percent efficiency. Instead, these open systems are shared by a large number of users, and the achieved application performance is often a small fraction of the peak hardware performance. This is not an agency-specific issue, but rather a shortfall in high-end computing capability that must be addressed by all agencies to serve their community's needs.

Achieving high-performance for complex applications requires a judicious match of computer architecture, system software and software development tools. Most researchers in high-end computing believe the key reasons for our current difficulties in achieving high performance on complex scientific applications can be traced to (a) inadequate research investment in software and (b) use of processor and memory architectures that are not well matched to scientific applications. Today, scientific applications are developed with software tools that are crude compared to those used in the commercial sector. Low-level programming, based on message-passing libraries, means that application developers must provide deep knowledge of application software behavior and its interaction with the underlying computing hardware. This is a tremendous intellectual burden that, unless rectified, will continue to limit the usability of high-end computing systems, restricting effective access to a small cadre of researchers.

Developing effective software (programming languages and tools, compilers, debuggers and performance tools) requires time and experience. Roughly twenty years elapsed from the time vector systems such as the Cray-1 first appeared in the 1970s until researchers and vendors developed compilers that could automatically generate software that operated as efficiently as that written by a human. This required multiple iterations of research, testing, product deployment and feedback before success was achieved.

In the 1990s, the U.S. HPCC program supported the development of several new computer systems. In retrospect, we did not learn the critical lesson of vector computing, namely the need for long-term, sustained and balanced investment in both hardware and software. We under-invested in software and expected innovative research approaches to high-level programming to yield robust, mature software in only 2–3 years. One need only look at the development history of Microsoft Windows™ to recognize the importance of an iterated cycle of development, deployment and feedback to develop an effective, widely used product. High quality research software is not cheap, it is labor intensive, and its successful creation requires the opportunity to incorporate the lessons learned from previous versions.

The second challenge for high-end computing is dependence on products derived too narrowly from the commercial computing market. Although this provides enormous financial leverage and rapid increases in peak processor performance, commercial and scientific computing workloads differ in one important and critical way—access to memory. Most commercial computer systems are designed to support applications that access a small fraction of a system's total memory during a given interval.

For commercial workloads, caches—small, high-speed memories attached to the processor—can hold the critical data for rapid access. In contrast, many, though not all, scientific applications (and several of those critical to signals intelligence and cryptanalysis) have irregular patterns of access to a large fraction of a system's memory. This is not a criticism of vendors, but rather a marketplace reality we must recognize and leverage. New high-end computing designs are needed to support these characteristics, both for fully custom high-end computer designs and more appropriate designs based on commodity components.

[1] According to the most recent "Top 500" list (www.top500.org), 6 teraflops at PSC, 6 teraflops at NCSA, 3.7 teraflops at SDSC and 10 teraflops at NERSC. Upcoming deployments at NCSA, SDSC and PSC will raise this number, but the aggregate will still be far less than user requirements.

The dramatic growth of the U.S. computing industry, with its concomitant economic benefits, has shifted the balance of influence on computing system design away from the government to the private sector. As the relative size of the high-end computing market has shrunk, we have not sustained the requisite levels of innovation and investment in high-end architecture and software needed for long-term U.S. competitiveness. Alternative strategies will be required.

This leads me to my third and final point: appropriate models of cooperation and support for high-end computing.

3. Actions

We must change the model for development, acquisition and deployment of high-end computing systems if the U.S. is to sustain the leadership needed for scientific discovery and national security in the long-term. The Japanese Earth System Simulator is a wakeup call, as it highlights the critical importance of both industry-government collaboration and long-term sustained investment. Reflecting the lessons of long-term investment I discussed earlier, the Earth System Simulator builds on twenty years of continued investment in a particular hardware and software model, and the lessons of six product generations. To sustain U.S. leadership in computational science, we must pursue two concurrent and mutually supporting paths, one short- to medium-term and the second long-term.

In the short- to medium-term, we must acquire and continue to deploy additional high-end systems at larger scale if we are to satisfy the unmet demand of the science and engineering research community. NSF's recent cyberinfrastructure report,[2] DOD's integrated high-end computing report, and DOE's ultrascale simulation studies have all made such recommendations. As one example, the cyberinfrastructure report noted "The United States academic research community should have access to the most powerful computers that can be built and operated in production mode at any point in time, rather than an order of magnitude less powerful, as has often been the case in the past decade." The cyberinfrastructure report estimated this deployment as costing roughly $75M/year per facility, with $50M/year per facility allocated to high-end computing hardware.

Given the interdependence between application characteristics and hardware architecture, this will require deployment of high-end systems based on diverse architectures, including large-scale message-based clusters, shared memory systems (SMPs) and vector systems.[3] Moreover, these systems must not be deployed in a vacuum, but rather must leverage another critical element of sustainable infrastructure—the experienced support staff members who work with application scientists to use high-end systems effectively. These high-end centers must also interoperate with a broad infrastructure of data archives, high-speed networks and scientific instruments.

High-end computing system deployments should not be viewed as an interagency competition, but rather as an unmet need that requires aggressive responses from multiple agencies. NSF and its academic supercomputing centers have successfully served the open academic research community for seventeen years; NSF should build on this success by deploying larger systems for open community access. Similarly, DOE has well served the high-end computing needs of laboratory researchers; it too should build on its successes. NIH, DOD, NASA, NSA and other agencies also require high-end capabilities in support of their missions, both for research and for national needs. The need is so large, and the shortfall is so great, that broader investment is needed by all agencies.

Concurrent with these deployments, we must begin a coordinated research and development effort to create high-end systems that are better matched to the characteristics of scientific applications. To be successful, these efforts must be coordinated across agencies in a much deeper and tighter way than in the past. This will require a broad, interagency program of basic research into computer architectures, system software, programming models, software tools and algorithms.

In addition, we must fund the design and construction of large-scale prototypes of next-generation high-end systems that includes balanced exploration of new hardware and software models, driven by scientific application requirements. Multiple, concurrent efforts will be required to reduce risk and to explore a sufficiently broad range of ideas; six efforts, each federally funded at a minimum level of $5M–$10M/year for five years, is the appropriate scale. At smaller scale, one will not be able

---

[2] *"Revolutionizing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure,"* January 2003, www.cise.nsf.gov/evnt/reports/toc.htm

[3] This is the approach we have adopted at NCSA, deploying multiple platforms, each targeting a distinct set of application needs.

to gain the requisite insights into the interplay of application needs, hardware capabilities, system software and programming models.

Such large-scale prototyping efforts will require the deep involvement and coordinated collaboration of vendors, national laboratories and centers, and academic researchers, with coordinated, multi-agency investment. After experimental assessment and community feedback, the most promising efforts should then transition to even larger scaling testing and vendor productization, and new prototyping efforts should be launched. It is also important to remember the lesson of the Earth System Simulator—the critical cycle of prototyping, assessment, and commercialization must be a long-term, sustaining investment, not a one time, crash program.

I believe we face both great opportunities and great challenges in high-end computing. Scientific discovery via computational science truly is the "endless frontier" of which Vannevar Bush spoke so eloquently in 1945. The challenges are for us to sustain the research, development and deployment of the high-end computing infrastructure needed to enable those discoveries.

In conclusion, Mr. Chairman, let me thank you for this committee's longstanding support for scientific discovery and innovation. Thank you very much for your time and attention. I would be pleased to answer any questions you might have.

Biography for Daniel A. Reed

Edward William and Jane Marr Gutgsell Professor; Director, National Center for Supercomputing Applications; Director, National Computational Science Alliance; Chief Architect, NSF ETF TeraGrid; University of Illinois at Urbana-Champaign

Dan Reed serves as Director of the National Computational Science Alliance (Alliance) and the National Center for Supercomputing Applications (NCSA) at the University of Illinois, Urbana-Champaign. In this dual directorship role, Reed provides strategic direction and leadership to the Alliance and NCSA and is the principal investigator for the Alliance cooperative agreement with the National Science Foundation.

Dr. Reed is one of two principal investigators and the Chief Architect for the NSF TeraGrid project to create a U.S. national infrastructure for Grid computing. The TeraGrid is a multi-year effort to build and deploy the world's largest, fastest, distributed computing infrastructure for open scientific research. Scientists will use the TeraGrid to make fundamental discoveries in fields as varied as biomedicine, global climate, and astrophysics. Dr. Reed is also the principal investigator and leader of NEESgrid, the system integration project for NSF's George E. Brown Jr. Network for Earthquake Engineering Simulation (NEES), which is integrating distributed instruments, computing systems, and collaboration infrastructure to transform earthquake engineering research.

Reed was head of the University of Illinois computer science department from 1996 to 2001 and, before becoming NCSA and Alliance director was co-leader of the Alliance's Enabling Technologies teams for three years. He is a member of several national collaborations, including the NSF Center for Grid Application Development Software, the Department of Energy (DOE) Scientific Discovery through Advanced Computing program, and the Los Alamos Computer Science Institute. He is chair of the NERSC Policy Board for Lawrence Berkeley National Laboratory, is co-chair of the Grid Physics Network Advisory Committee and is a member of the board of directors of the Computing Research Association (CRA). He recently served as program chair for the workshop on the Road Map for the Revitalization of High End Computing for Information Technology Research and Development (NTTRD). He is an incoming member of the President's Information Technology Advisory Committee (PITAC). In addition, he served as a member of Illinois Governor's VentureTECH committee, which advised the former governor on technology investment in Illinois.

UNIVERSITY OF ILLINOIS
AT URBANA-CHAMPAIGN

National Center for Supercomputing Applications
152 Computing Applications Building
605 East Springfield Avenue
Champaign, IL 61820

July 14, 2003

The Honorable Sherwood Boehlert
Chairman, Committee on Science
2320 Rayburn Office Building
Washington, DC 20515

Dear Congressman Boehlert:

Thank you for the invitation to testify before the U.S. House of Representatives Committee on Science, on July 16, 2003 for the hearing entitled "Supercomputing: Is the U.S. on the Right Path?" In accordance with the Rules Governing Testimony, this letter serves as formal notice of the Federal funding I currently receive in support of my research.

Please see attached table containing federal awards over the past three years.

Sincerely,

Daniel A. Reed

Daniel A. Reed
Edward William and Jane Marr Gutgsell Professor
Director, National Center for Supercomputing Applications
Director, National Computational Science Alliance
Chief Architect, NSF TeraGrid

Attachment

telephone 217 244-0072

Grant Number	Federal Agency/Source	Title	Project Dates	Total Project Amount
N00014-03-1-0765	Office of Naval Research	National Center for Advanced Secure Systems Research	05/05/2003 - 04/27/2004	4,195,814
DE-FC02-01ER41205	Department of Energy	National Computational Infrastructure for Lattice Gauge Theory	08/15/2001 - 08/14/2004	$ 273,273
ARMY RAYTHEON AA19	Raytheon Company	NCSA Participation in High-Performance Computing Modernization Program at Army	09/23/1996 - 08/17/2004	3,278,490
NSF ACI-0219597	National Science Foundation	ITR: Intelligent High-Performance Computing on Toys	09/01/2002 - 08/31/2004	$ 400,000
NSF EIA 9972884 EQ	National Science Foundation	Intelligent Information Spaces: A Testbed to Explore and Evaluate	09/15/1999 - 08/31/2004	$ 2,000,000
DE-FC02-01ER25488	Department of Energy	High-End Computer System Performance: Science and Engineering	09/15/2001 - 09/14/2004	$ 801,000
NSF CMS 01-17853	National Science Foundation	NEESGRID: A Distributed Virtual Laboratory for Advanced Earthquake Experiment	08/01/2001 - 09/30/2004	6,501,869
NSF ACI 01-22296	National Science Foundation	The TERAGRID: Cyberinfrastructure for 21ST Century Science and Engineering	10/01/2001- 9/30/2004	36,045,500
NSF EOA 9975248	National Science Foundation	CADRE: A National Facility for High-Performance I/O Characterization	07/15/1999 - 06/30/2003	$ 1,200,000
R71700H-29200099	William Marsh Rice University (Department of Energy Prime)	Los Alamos Computer Science Institute	05/03/1999 - 09/20/2003	$ 1,164,909
R38144-79200003	William Marsh Rice University (National Science Foundation Prime)	NGS: GrADS: Efficient Script-Based Application Development for Networked High Performance Computing Environment	10/01/2002 - 09/20/2003	$ 105,000
NSF ACI-9619019	National Science Foundation	Enabling Technology Team C: I/O and Virtual Environments Software	10/01/1997 - 09/30/2003	$ 800,000
NSF ACI 96-19019COOP	National Science Foundation	National Computational Science Alliance - PACI	10/01/1997 - 9/30/2003	192,846,587
DOE LLNL B505214	Lawrence Livermore National Lab (Department of Energy Prime)	Wide-Area, Adpative I/O Systems for Data and Visualization Corridors	07/29/1999 - 07/20/2002	$ 702,819
PC253664	California Institute of Technology (Department of Energy Prime)	Accelerated Strategic Computing Initiative	10/01/1997 - 11/30/2002	$ 515,658
R36504-29200099	William Marsh Rice University (National Science Foundation Prime)	Grid Appplication Development Software (GrADS)	10/01/1999 - 12/31/2002	$ 445,935

Chairman BOEHLERT. Thank you, Dr. Reed. And I might say you lived up to the advanced billing so ably presented by Mr. Johnson.

Dr. REED. Thank you.

Chairman BOEHLERT. Mr. Scarafino.

STATEMENT OF MR. VINCENT F. SCARAFINO, MANAGER, NUMERICALLY INTENSIVE COMPUTING, FORD MOTOR COMPANY

Mr. SCARAFINO. Thank you for the opportunity to testify regarding the national needs for advanced scientific computing and industrial applications. My name is Vincent Scarafino, and I am Manager of Numerically Intensive Computing for Ford Motor Company.

Ford has a long and proud history of leadership in advancing engineering applications and technologies that stretches back over the last 100 years. Ford uses high-performance computing to help predict the behavior of its products in nominal and extreme conditions. Approximately half of the computing capacity is used for safety analysis to determine what its vehicles will do in impact situations. Other uses include vehicle ride performance, usually referred to as NVH, noise, vibration, and harshness, which is predicted using model analysis techniques. Fluid flow analysis is used to predict such things as drag characteristics of vehicles, under hood temperatures, defroster and interior climate control, catalytic converter design, and exhaust systems.

Computing capabilities allow Ford to accelerate the design cycle. Better high-end computing systems will help engineers balance competing design requirements. Performance, durability, crash worthiness, occupant and pedestrian protection are among them. These tools are necessary to stay in business. Competition from national and international companies is intense.

Significantly faster high-end machines would help improve the ability to predict vehicle safety as well as the durability and wind noise characteristics, which are among the most cited customer issues. Although safety codes now run well on—in parallel on commodity-based clusters, high-end machines would provide improved turnaround times, leading to reduced design cycle time and the opportunity to explore a greater variance of design parameters. Advanced durability and wind noise analyses typically require computer runs greater than two weeks, and this is with coarse models on the fastest machines we have access to. The durability work runs on 4-way SMP 1.25 GHz Alpha processors. The wind noise runs with 6-way SMP on a Cray T90. They cannot effectively use more processors.

Up until the mid-1990's, the Federal Government had helped with the development of high-end machines with faster, more powerful processing capability and matching memory bandwidth and latency characteristics by helping to fund development and create a market for them. These machines were built mainly to meet the needs of government security and scientific research. Once they were built, there was a limited, but significant application of these machines in the private sector. The availability of higher capability machines advanced the application of science in the private sector. This worked quite well.

In the mid-1990's, the Federal Government decided to rely on utilizing off-the-shelf components and depend on the ability to combine thousands of these components to work in harmony to meet its advanced high-performance computing needs. The result was an advance in the areas of computer science that dealt with parallel processing. Over the last eight years, some kinds of applications have adapted well to the more constrained environment supported by these commodity-based machines. Vehicle safety analysis programs are an example. Most vehicle impact analysis can now be done on commodity clusters. We have learned how to do "old science" considerably cheaper.

We have not made any significant advancement in new science, some examples of which would include advanced occupant injury analysis and the modeling of new complex materials, such as composites. The physics is more complex, and the computational requirements are beyond current capability. The hardest problems do not adapt well to parallel architectures. Either we don't know enough about the problem to develop a parallel solution, or they do not—they are not parallel by nature.

The Federal Government cannot rely on fundamental economic forces to advance high-performance computing capability. The only economic model with enough volume to drive this kind of development is the video game industry. Unfortunately, these models—these machines are very fast only for a particular application and do not provide a general solution. The Federal Government should help with the advancement of high-end processor design and other fundamental components necessary to develop well-balanced, highly capable machines. U.S. leadership is currently at risk.

The advanced computing needs of the scientific community and the private sector are not diverging. The fact that there has been no fundamental advance in high-performance capability computers in the last eight years has forced these communities to adapt less qualified commercial offerings to the solution of their problems. If advanced computing capability becomes available in the form of next-generation supercomputers, the scientific community and the private sector would be able to utilize them for the application of new science.

Thank you.

[The prepared statement of Mr. Scarafino follows:]

PREPARED STATEMENT OF VINCENT F. SCARAFINO

Thank you for the opportunity to testify regarding the national needs for advanced scientific computing and industrial applications. My name is Vincent Scarafino and I am Manager of Numerically Intensive Computing for Ford Motor Company. Our automotive brands include Ford, Lincoln, Mercury, Volvo, Jaguar, Land Rover, Aston Martin and Mazda.

Ford has a long and proud history of leadership in advanced engineering applications and technologies that stretches back over the last 100 years. Henry Ford set the tone for the research and development that has brought many firsts to the mass public. In the future, the company will continue to apply innovative technology to its core business. New technologies, such as hybrid electric and fuel cell powered vehicles, will help achieve new goals, reward our shareholders and benefit society. We understand that the quality and safety or our products are fundamental to our corporate success.

For example, to be among the leaders in vehicles safety, we must commit ourselves to ongoing improvement of the safety and value our products. We do this through research, product development and extensive testing of our products.

Ford uses high-performance computing to help predict the behavior of its products in nominal and extreme conditions. Approximately half of the computing capacity is used for safety analysis to determine what its vehicles will do in impact situations. Other uses include vehicle ride performance (usually referred to as NVH—noise, vibration and harshness), which is predicted using modal analysis techniques. Fluid flow analysis is used to predict such things as drag characteristics of vehicles, under hood temperatures, defroster and interior climate control, catalytic converter design, and exhaust systems.

Computing capabilities allow Ford to accelerate the design cycle. Better high-end computing systems will help engineers balance competing design requirements—performance, durability, crash-worthiness, occupant and pedestrian protection are among them. These tools are necessary to stay in business. Competition from national and international companies is intense.

Significantly faster high-end machines would help improve the ability to predict vehicle safety as well as durability and wind noise characteristics, which are among the most cited customer issues. Although safety codes now run well in parallel on commodity based clusters, high-end machines would provide improved turnaround times leading to reduced design cycle time and the opportunity to explore a greater variance of design parameters. Advanced durability and wind noise analyses typically require computer runs greater than two weeks, and this is with coarse models on the fastest machines we have access to. The durability work runs on 4-way SMP 1.25 GHz Alpha processors. The wind noise runs with 6-way SMP on a Cray T90. They cannot effectively use more processors.

Up until the mid 1990's, the Federal Government had helped with the development of high-end machines with faster, more powerful processing capability and matching memory bandwidth and latency characteristics by helping to fund development and create a market for them. These machines were built mainly to meet the needs of government security and scientific research. Once they were built, there was a limited, but significant application of these machines in the private sector. The availability of higher capability machines advanced the application of science in the private sector. This worked quite well.

In the mid 1990's the Federal Government decided to rely on utilizing off-the-shelf components and depend on the ability to combine thousands of these components to work in harmony to meet its advanced high-performance computing needs. The result was an advance in the areas of computer science that dealt with parallel processing. Over the last eight years, some kinds of applications have adapted well to the more constrained environment supported by these commodity based machines. Vehicle safety analysis programs are an example. Most vehicle impact analysis can now be done on commodity clusters. We have learned how to do "old science" considerably cheaper. We have not made any significant advancement in new science. Some examples include advanced occupant injury analysis and the modeling of new complex materials such as composites. The physics is more complex and the computational requirements are beyond current capability. The hardest problems do not adapt well to parallel architectures. Either we don't know enough about the problem to develop a parallel solution, or they are not parallel by nature.

The Federal Government cannot rely on fundamental economic forces to advance high-performance computing capability. The only economic model with enough volume to drive this kind of development is the video game industry. Unfortunately these machines are very fast only for a particular application and do not provide a general solution. The Federal Government should help with the advancement of high-end processor design and other fundamental components necessary to develop well-balanced, highly capable machines. U.S. leadership is currently at risk.

The advanced computing needs of the scientific community and the private sector are not diverging. The fact that there has been no fundamental advance in high-performance capability computers in the last eight years has forced these communities to adapt less qualified commercial offerings to the solution of their problems. If advanced computing capability becomes available in the form of next generation supercomputers, the scientific community and the private sector would be able to utilize them for the application of new science.

Ford continues to invest time, money and a significant portion of our company's human resources to study and explore how our vehicles can perform even better for the next 100 years. We believe that the next generation of supercomputers is essential to ensure that our industry, which has contributed greatly to economic growth in this country, remains competitive.

BIOGRAPHY FOR VINCENT F. SCARAFINO

Vince Scarafino is currently Manager of Numerically Intensive Computing at Ford Motor Company. He has been with Ford for close to 30 years and has over 35 years of experience with computers and computing systems. Most of his work has been in the engineering and scientific areas. He started with GE 635, IBM 7090, and IBM 1130 machines. His experience with security and relational data base technology on Multics systems lead to sharing knowledge with academic and government organizations. His efforts in the last decade have been dedicated to providing flexible and reliable supercomputer resources for the vehicle product development community at Ford via an open systems-based grid environment connecting three data centers on two continents. He currently manages a broad range of high performance computers, from Cray T90 to Linux Beowulf Clusters, that are used by Ford's engineers around the world for such things as vehicle design and safety analysis.

Vince has a degree in mechanical engineering from the University of Michigan.

September 29, 2003

The Honorable Sherwood Boehlert
Chairman, Science Committee
2320 Rayburn Office Building
Washington, DC 20515

Dear Congressman Boehlert:

Thank you for the invitation to testify before the U.S. House of Representatives Science Committee on July 16 for the hearing entitled *Advanced Scientific Computing: Are We on the Right Path?* In accordance with the Rules Governing Testimony, this letter serves as formal notice of the Federal funding Ford Motor Company currently receives in support of my research.

- Ford Motor Company received no federal funding directly supporting the subject matter on which I testified, in the current fiscal year or either of the two proceeding fiscal years.

Sincerely,

Vincent F. Scarafino

DISCUSSION

Chairman BOEHLERT. Thank you very much.

Let me start by asking Dr. Orbach and Dr. Freeman. This interagency coordination that we think is so vital, as I understand it, that is sort of being shepherded by OSTP. And how long has that been in operation, this interagency coordinating vehicle? What are they called? Council or—Dr. Orbach?

Dr. ORBACH. It is the High-End Computing Revitalization Task Force. That Committee is expected to finish its report in August, next month, and has been in progress for about six months.

Dr. FREEMAN. A bit less, I believe. It started in the April time frame, I believe.

Chairman BOEHLERT. Okay. But the report is not the end of the job. I mean, they are going to—it is going to continue, isn't it?

Dr. FREEMAN. I am not certain of the intentions of OSTP.

Chairman BOEHLERT. Do you think it should continue? In other words, don't you think we need an interagency coordinating vehicle to carry forward on a continuing basis a program where the principals of the agencies involved get together, not to just chat and pass the time of day, but get really serious about a really serious subject?

Dr. FREEMAN. I certainly agree that coordination is necessary. Let me remind you that there has been in existence, for over 10 years, the National Coordination Office for Networking and IT R&D. Indeed, that office, the National Coordination Office, is managing the High-End Computing Revitalization Task Force that we were just discussing.

Chairman BOEHLERT. That has been in existence for a decade at a time when—instead of going like this, as we all want to, and the recorder can't say "like this", so let me point out a dramatic upward turn, it sort of leveled off. And we are getting the pants beat off of us in some areas.

Dr. FREEMAN. I think my point, Mr. Chairman, is that the mechanism is in place for coordination. As you know, I have only been in the government for about a year, so I can't speak to the total history of that office, but I think, certainly—and I chair or co-chair the Interagency Working Group. But the—what I hear you calling for is a higher level coordination. I think some steps have already been taken in that direction that Dr. Orbach may be able to speak to.

Chairman BOEHLERT. But would you agree that OSTP is the place, maybe, to be the coordinating vehicle for this interagency cooperative effort? And don't be turf-conscious. I mean, just sort of give it some thought. I mean, the Science Advisor to the President, head of the Office of Science and Technology Policy, it would seem to me that would be the ideal vehicle and that the principals not send a designee to that meeting, or at least every meeting. Once in a while the principals, including Dr. Marberger and Dr. Freeman and Dr. Orbach, you know, you get there, not some designee. And you really talk turkey.

And then it seems to me that as you go in your budget preparation for the next years, you are going into OMB. You go so—forward with a coordinated plan, so that if OMB is looking at

Orbach's recommendation for DOE, they will be mindful of what Freeman's recommendation is for NSF and what somebody else's is for DARPA, in other words, a coordinated approach. This is too big a challenge to not just really put our shoulder to the wheel, so to speak, and that is sort of an uncomfortable position, but you know exactly what I am saying.

Dr. Orbach, do you have any comment on that?

Dr. ORBACH. Mr. Chairman, I subscribe to your remarks directly. In fact, the principals did meet, under Dr. Marberger's direction, a few weeks ago to discuss precisely the issues, which you laid out.

Chairman BOEHLERT. And they are going forward?

Dr. ORBACH. And they are going forward. And my personal preference would be that OSTP would continue that relationship and leadership. When the principals met, there was unanimity that it was—needed to be a collaborative affair where each of us brought our areas of expertise to the table and work together.

Chairman BOEHLERT. And the coordinated approach to the budget process is extremely important. Speaking about the budget process, what about beyond '04, Dr. Freeman, with NSF's plans for the Supercomputing Centers? Where do you envision going in '05 and '06?

Dr. FREEMAN. Well, as I noted in my both verbal and written testimony, although we have announced plans to change the modality of funding in line with the recommendations of the Atkins Report, we intend that the amount of money going into the activities that are currently covered by that PACI program, in fact, will increase. We have announced some aspects of that. As I noted, other aspects are currently under internal preparation and will be announced and vetted through our National Science Board in the coming month.

Chairman BOEHLERT. So is that in the future I should be optimistic about our centers beyond '04?

Dr. FREEMAN. Absolutely, and I——

Chairman BOEHLERT. And I should be optimistic that resources will come to do some of the things that need to be done in terms of purchasing hardware, for example?

Dr. FREEMAN. That is certainly our intention.

Chairman BOEHLERT. Thank you very much. And Mr. Johnson, if he were here, would be comforted by that answer.

Let me ask——

Dr. FREEMAN. As well as Dr. Reed, I might add.

Chairman BOEHLERT. Well, I know. I didn't want to be too obvious.

Let me turn to our non-government witnesses, Dr. Reed and Mr. Scarafino. And I thank you for your input. Is the Federal Government policy, as you understand it, and I am not quite sure it is crystal clear, do you think we are moving in the right direction? And if not, where would you put some additional focus? We will start with you, Dr. Reed.

Dr. REED. Well, I agree, absolutely, that I think greater coordination is needed. I think one of the challenges for us—as I said, there are really two aspects of this problem. There is what we do in the short-term and the medium-term to address the shortfall of the scientific needs, vis-á-vis the Earth Simulator and what is available

to the science community. And there is a clear shortfall. And I realize everyone that comes before you pleads poverty and asks for more money. I mean, that is—I know that. But there is a clear need for additional resources to address the community. I think it is fair to say that in terms of one sort of breakthrough calculation, we, perhaps, have, in the open community, enough supercomputing capability for, perhaps, solving one problem, at the moment. And there are many problems at that scale that the community would like to solve. So that is one of the issues.

The longer-term issue, and the place where I think coordination is equally critical in addressing the issue that Mr. Scarafino addressed as well, and that is the commercial application of technology to high-performance computing. There is a disconnect in the capability that we can procure, regardless of the amount of money we have, to address certain needs. And some of those machines simply aren't available at any price right now.

And the place where coordination would be enormously helpful is not only—if you think about the research pipeline from the basic long-term research that is required to investigate the new architectures and new software and new algorithms for applications, the place where we have had a gap, in some sense, has been the building of machines at sufficiently large scale to test them so that industry would pick them up and commercialize them and make them products. That gap is a place where coordination across DARPA, DOE, NSF, and other agencies could really play a critical role. And I think one of the things we have to do there is fund coordinated teams across those agencies to address some of those challenges. That is a place where, operationally, a very big difference could be made. And that was a message that was echoed very loudly at the workshop I chaired, which was the community input workshop to the process that Dr. Freeman and Dr. Orbach discussed where the science and engineering community was providing responses to what the tentative agency plans were.

Chairman BOEHLERT. Mr. Scarafino?

Mr. SCARAFINO. I am not sure exactly what kind of mechanisms could be best used to accomplish these things, but it is very good that the Federal Government is awakened to the shortcomings in the very-high computing arena and that they are basically taking steps to correct that. That is——

Chairman BOEHLERT. So I would imagine within Ford—you notice I gave you a plug in my introductory statement. I said when the light bulb went on, we had a better idea. Thank you very much for that response.

The Chair recognizes Mr. Bell.

Mr. BELL. Thank you, Mr. Chairman. And thank you for calling this hearing. I would agree with the Chair and Ranking Member that it is an extremely important topic, and I appreciate your testimony here today.

I want to follow up on an issue that Chairman Boehlert raised in his opening statement regarding Japan and the fact that they are recognized as the world leader in supercomputing and certainly have a lead on us and combine that with another subject that this committee has paid a great deal of attention to over the course of this year, and that is nanotechnology, because in a similar parallel,

Japan is also a world leader in the area of nanotech. As you all know, in the future, it is likely that nanotech particles or a nanochip will be used to replace the silicon chip, and nanotech products or nanochips are smaller and can hold more memory and have a lot of advantages. And since it is the role of the government to continue to fund research on new approaches to building high-performance computing hardware or supercomputers and nanotech seems to be a means to this end, can you tell me, and perhaps beginning with Dr. Orbach and anybody else who would like to comment, in what ways the government is looking at nanotechnology to create or develop new supercomputing hardware and any progress that is being made in that area?

Dr. ORBACH. It works both ways. Not only can nanotechnology assist in the creation of new architectures, but also the high-end computing is required to understand the properties of nanoscale materials. The Office of Science had a workshop on theory and modeling in nanoscience, and discovered that computational capabilities, exactly as Mr. Scarafino talked about with regard to composites, is essential to really understand what it is you are doing when you are creating these new materials at the atomic level. So it works both ways and is a critical need.

Mr. BELL. Dr. Reed, I saw you—or Dr. Freeman, did you want to comment?

Dr. FREEMAN. Yes, I would just note that NSF is heavily involved in the government-wide nanotechnology initiative, as you know. One of the missions of the directorate that I head is to explore, if I may use the vernacular, far out possibilities for computation: quantum computing, biological computing in the sense of using biological materials to compute. And I am sure that some of our investigators are looking at some of the types of applications that you referred to as well.

Mr. BELL. Dr. Reed?

Dr. REED. I just want to echo what was said. It is certainly the case that applications of high-end computing are really critical to exploring novel materials. One of the things that we are close to being able to do, that we can see on the horizon, is looking at ab initio calculations from first principles to design new composite materials. And that is one of the applications.

To hop back to the lead-in of your question about the Earth System Simulator and that aspect of it, I think one of the lessons to draw from that effort is the importance of long-term investment. The Earth Simulator was not an overnight sensation. It was a long period of investment. There were roughly 20 years of R&D behind that machine and six generations of precursor products where ideas and technologies were combined to build that. And that is why earlier in my oral testimony I said that sustained investment is really part of the key to developing these technologies.

But Dr. Freeman and Dr. Orbach are absolutely correct that there is an interplay. It goes both ways. The application of high-end computing can design materials, and new materials are really critical in designing the next generation of machines.

Mr. BELL. Mr. Scarafino?

Mr. SCARAFINO. I don't really have any expertise in this area, so there is nothing I could add.

Mr. BELL. Fair enough. And Dr. Reed, maybe I should ask you, or any of you that know, what Japan is doing in regard to nanotechnology as it relates to supercomputing and what the U.S. needs to do, in your opinion, to catch up. Do you know?

Dr. REED. I can't comment specifically on the nanotechnology aspect. I can say that the Japanese have historically had long-term investment in building a whole series of machines of which the Earth Simulator is just one example. One example of a machine that they have built at least six generations of, as well, is focused on machines designed specifically to support high-energy physics. There have been a series of machines called GRAPE, up through, I believe, GRAPE–6, that are focused on the same kinds of problems that we looked at here in trying to compute from first principals the masses of elementary particles. And those machines have been designed working directly with the physicists who want to solve those problems to understand the match of the application requirements to the architecture and the software.

And that interplay of the scientific computing needs with the design on the hardware and the software is where the—often a disconnect has been when we have tried to leverage commercial hardware, because it is driven, as Mr. Scarafino mentioned, by—primarily by the needs of the commercial market. And that interplay is one of the things that I think has been the hallmark of successful scientific computing over the years at the National Laboratories, at the centers. It has been looking at the way hardware, architecture, software, driven by applications, can shape the future.

Dr. FREEMAN. I would certainly agree with Dr. Reed. But I would just add that I have a meeting this afternoon, indeed, to help shape a study that we intend to kick off in cooperation with several other agencies, specifically to make sure that we are well advised as to what is going on in Japan in some of these forefront areas.

Mr. BELL. And then we would have a better idea of what would be required to——

Dr. FREEMAN. Precisely.

Mr. BELL. Dr. Orbach, any thoughts on that?

Dr. ORBACH. I would second your observation. The Japanese have been very aggressive in the nanotechnology area. And we will provide, for the record, some of the details to bring you up—all of us up to speed on that, if that would be acceptable.

[Note: This information is located in Appendix 2: Additional Material for the Record.]

Mr. BELL. It would be, and I appreciate it.

And thank you, Mr. Chairman.

Chairman BOEHLERT. Thank you, Mr. Bell.

And now it is a pleasure to turn to a distinguished scientist and educator, a fellow of the American Physical Society, the gentleman from Michigan, Dr. Ehlers.

Mr. EHLERS. Thank you, Mr. Chairman. That is almost as glowing as Dr. Reed's introduction.

I am fascinated listening to this and particularly the discussions of teraflops, et cetera. I go back to the days when I was happy to get one flop. I cut my eyeteeth on the IBM–650, which I believe was developed before Dr. Reed was even born. And I was delighted

when—and I was one of the first to use that in my field of physics and found it extremely useful.

I—two questions. The first one, I am pleased with the memorandum of understanding that you referenced, Dr. Orbach, in working with other agencies and the cooperation that is being discussed between DOE and NNSA, both your branch and the NNSA, as well as DARPA and NSF. But what about some of the other agencies, for example, NOAA [National Oceanic and Atmospheric Administration] and its affiliate NCAR [National Center for Atmospheric Research], use—have a great need for very high-performance computers. How are they involved in all that is going on, and how do you share resources with them? Are they part of this task force? Are they part of the decision-making process? Because clearly, climate modeling is of an immense interest, particularly in relationship to climate change, but even in the mundane aspects of everyday weather projection. Dr. Orbach, do you want to kick that off?

Dr. ORBACH. Yes, we are working—the—very closely, the Department of Energy Office of Science has a responsibility for climate modeling among all of the other agencies. And we are working very closely with NOAA and the other programs to create computational architectures, which will help in modeling. It is interesting to note that the Earth Simulator can do much better predictions, long-term predictions, on climate than we can by almost a factor of 10. Their grid size is that much smaller. So it is just evidence of the computational power that these machines can provide.

Mr. EHLERS. Dr. Freeman, anything to add?

Dr. FREEMAN. Let me just note, sir, that I believe you mentioned NCAR. That is actually a federally-funded research and development center that is largely supported by the National Science Foundation. And so through that mechanism, they are very intricately involved with our plans at NSF. Likewise, in terms of the OSTP planning activity, I believe NOAA is involved there, as are, essentially, all of the agencies that you might expect to be engaged in this activity.

Mr. EHLERS. Thank you.

My second question, and that has to do with, Dr. Orbach, with DOE's plans to provide supercomputer resources to academic researchers. And I am wondering how this is going to relate to the NSF Supercomputer Centers. Are—is there going to be a meshing here, or is this simply going to provide another avenue for academic researchers? They can knock on the door of a Supercomputing Center. If they don't get what they need, they knock on the door of DOE and say, "What do you have to offer?" Could you amplify on how that is—how that interrelationship is going to work out?

Dr. ORBACH. Well, the interrelationship is very important, because the National Science Foundation grid computation structure will enable researchers everywhere in the United States to couple into the high-end computational program. So we rely on NSF for those relationships. We are going to be opening our largest computer, NERSC, at Lawrence Berkeley National Laboratory to grand challenge calculations. Ten percent of NERSC will be available to any scientist, regardless of where their support comes from, actually internationally, to compete for a grand challenge calculation.

The opportunity to couple depends on the network of the cyberinfrastructure, which you have heard referred to before, and so we are working in a very interdependent way to enable access to these machines.

Mr. EHLERS. Thank you.

My final question and concern is, first of all, what is going on now is wonderful, but I think it is very late. A number of years ago, the Japanese took the lead. We have maintained our present more by a force of the Federal Government saying, "We are only going to buy American computers rather than Japanese," even though we did get some Japanese ones. But we lost our edge. How did that happen, and is this really going to help us regain our edge in the international arena? I will open it up to—Dr. Reed, you are itching to answer this one.

Dr. REED. I don't know if I am itching to answer, but I will try.

I think one of the things that happened was that what happened, if you look historically in—to hop back, perhaps not to the 650, but certainly to earlier days, the machines that were designed during the height of the Cold War were driven clearly by national needs for security and weapons design. And there was a compelling national need as well as a market opportunity. As the computing market overall grew, high-performance computing became a smaller fraction of the overall market, and therefore the financial leverage that the Federal Government had at the high-end has declined some. It is still substantial and dominant, but it has declined. And that has led manufacturers to use commercial products, or variations thereof, to try to address those needs.

But I think the critical thing that happened was there is a gap between basic research and next-generation architecture to the single investigator model and production. And what happens in that intermediate phase is testing ideas at sufficient scale that vendors see that a productization of that idea is possible and viable and likely to succeed. It is much less common for ideas from papers to have a revolutionary effect on new-generation designs. They have an evolutionary effect, and we have seen that in commercial microprocessor design. But something at sufficiently prototype scale that you can see with real applications that there are benefits that you can get the fractions of peak performance that machines, like the Earth Simulator, have demonstrated.

That gap of testing those ideas at scale to provide the push to get it over the energy barrier, if you will, to productization is one of the things that we haven't had. And I think if we could rekindle that effort, and it did very much go on in the 1990's, even, with the HPCC [High Performance Computing and Communications] Program, and before that, in the '80's and '70's, we would go a long way to regaining the lead.

As I said, it is not a one-time crash effort. It is a process that really needs to continue and be sustained at some level to feed new ideas into the process.

Mr. EHLERS. And how did the Japanese get their edge?

Dr. REED. They sustained an investment in machines. As we looked at vector machines, the conventional wisdom, if you will, was that that was not a path to continue. And in the early 1990's, we decided, and I use "we" in the broad sense, decided that there

was not much more headroom in that direction, and we started to explore other alternatives, which led to the parallel computing market that Mr. Scarafino referred to earlier. But we largely stopped development in those vector machines that capitalized on very efficient access to memory that supported scientific applications. Almost no development happened there and little investment for a period of 10 years or more.

The Japanese continued to invest in those machines to build next-generation products. And that is why I said the Earth Simulator is really the product of 20 years investment and at least a half a dozen product generations. We went in a different direction, which yielded some real benefits, but we stopped going in a direction that still had additional opportunity.

Mr. EHLERS. Thank you.

Chairman BOEHLERT. Thank you.

Mr. Udall.

Mr. UDALL. Thank you, Mr. Chairman. I want to thank the panel as well and apologize. I was at another hearing, and I hope I am not redundant in my question. But I do want to pick on my colleague, Mr. Ehlers' questions about parallel and vector computing and direct the question to Dr. Freeman and Dr. Orbach.

Are there fields in the area of science and engineering where progress has been delayed because of inadequate access to these high-end computing technologies? I am mindful of the Clinton Administration when they did a national climate assessment study as—in the last years of that Administration, I believe—had to ask the indulgence of the Europeans and the Canadians and the British for computing power to actually draw that assessment on climate. That is an interesting dynamic, to say the least. Would you all comment on my question and—if you would?

Dr. FREEMAN. I think, certainly, some of the testimony this morning and certainly it is the case that there are very important problems that no one, including the Japanese, can properly solve today. And there are certainly those that—where our progress or our particular level of expertise may be retarded because of our not having the, not only capability, but the capacity, that is enough supercomputers to go around, if you will.

Mr. UDALL. Are there particular fields, Dr. Freeman or Dr. Orbach that——

Dr. ORBACH. Yes, I wanted to amplify on that. Indeed, there are fields.

Mr. UDALL. Um-hum.

Dr. ORBACH. There are—we have had about eight virtual workshops from across the scientific spectrum that have looked at opportunities that high-end computation could provide. But I would like to turn it around a little bit. I think the Japanese are to be congratulated for construction of this machine. What it did is it showed us that it was possible to get very high sustained speeds on problems of scientific interest. They were able to get 26.5 teraflops on geophysics problems. And as you heard previously, we think that the range of somewhere between 25 to, maybe, 50 teraflops opens up a whole new set of opportunities——

Mr. UDALL. Um-hum.

Dr. ORBACH [continuing]. For science that had never been realized before. And so what we have, thanks to the Japanese now, is an existence proof, namely that it is possible to generate these high sustained speeds on specific scientific questions. Now whether it will be one or two or three different architectures, which will be required to address these, is what we are investigating now. So the United States is pursuing assiduously exploration of different architectures to see which is best suited for the problems at hand.

Mr. UDALL. Um-hum. Dr. Reed, you seem to be energized by the question as well. Would you like to comment?

Dr. REED. Well, there is a long list of application domains. One of the things that came out at the workshop that I mentioned earlier, as well as other workshops, has been talking to application scientists about the opportunities that they see. A couple of examples as illustration, one that is of intellectual interest and one that is of very practical interest. Intellectual interest, there have been huge new results from observational astronomy that suggest that most of the matter and energy that we see is, in fact, a small fraction of the universe. And so trying to understand dark matter and dark energy and build models of the large scale evolution of the universe, in essence, to answer one of the really big questions: how did everything get here?

Mr. UDALL. Um-hum.

Dr. REED. It is a problem that suddenly becomes possible. And it is a compelling example of the power of high-performance computing, because we can't do experiments on the universe. We only have one. It is what it is. But we can evolve models of it from first principles and compare those with observational data. That is really one of the unique capabilities of computational modeling, in comparison to experiment and theory.

Another one, of course, is looking forward to rational drug design and being able to understand how proteins fold and how various molecules can dock with those. Those are problems that are out of range right now, but with additional computing capability, have very practical benefits, because a huge fraction of the expense that pharmaceutical companies incur is in assessing potential drug candidates. They discard thousands for every one that even reaches first stage clinical trials. And high-performance computing can compress that time scale. It can reduce the costs for them.

Mr. UDALL. Who would own that computer that would be utilized by the pharmaceutical companies?

Dr. REED. Well, that is an interesting history there, because it is a great example of the interplay between government investment and industry. In fact, one of the very first technology transfers that my center, NCSA, was engaged in years ago, was where we transferred a lot of early vector computing technology. They actually bought their own Cray vector supercomputer and used for some early drug designs. So there were connections and collaborations where that technology can flow.

Mr. UDALL. It sounds like a nice example of the public/private partnership and public dollars going to public but also private benefit. Thank you.

Chairman BOEHLERT. Thank you very much.

Mr. UDALL. Thank you, Mr. Chairman.

Chairman BOEHLERT. The gentleman's time is expired.

It is a pleasure to welcome the newest Member of our Committee, and a valued addition to the Committee, I might add, Mr. Randy Neugebauer from Texas.

Mr. NEUGEBAUER. Thank you, Mr. Chairman.

Mr. Reed, do you believe the NSF is providing adequate support for the Supercomputer Centers? And have you detected any change in policy with regard to NSF's commitment to the Supercomputer Centers?

Dr. REED. Well, as I said at the outset, NSF has been a long and steadfast supporter of high-performance computing, yet the centers have a 17-year history, as Chairman Boehlert mentioned before. They were begun in the 1980's. And we look forward to continued investment. We also recognize that NSF has a portfolio of investments in infrastructure and research. And budgets are constrained at the moment.

Having said that, we do certainly see untapped opportunities where additional capability and investment, not only in the hardware, but in the people and the support infrastructure, could attack some new problems.

Mr. NEUGEBAUER. Thank you.

Chairman BOEHLERT. Boy, was that a diplomatic answer.

Mr. NEUGEBAUER. Exactly.

Dr. Reed, and Mr. Scarafino, Dr. Reed's testimony states that there is a large and unmet demand for access to high-end computing in support of basic scientific and engineering research. Dr. Freeman and Dr. Orbach have described their agencies' plans for the future. And this question could be to Dr. Reed and to Mr. Scarafino. Do you believe that these plans are adequate to provide more high-end computing systems and then enable the development of the new capabilities?

Dr. REED. Why—as I said before, and as you quoted from my written testimony, I do think there is an unmet demand for capability. At a high-level, I think it is fair to say that we can perhaps solve one or two of these really critical problems. As Mr. Bell asked about nanotechnology, a large-scale effort to understand new nanomaterials itself could largely consume the available open capacity that we have at the moment.

There are a whole set of problems like that, and it is a matter of making priority choices about what the level of investment is and where one will accelerate progress. There is no doubt that additional investment would yield new progress, however.

Mr. SCARAFINO. Since about the mid-1990's, we have kind of been at a standstill with regard to the availability of more capable computers. The area where we are struggling with as far as making advances is in the area of durability and in the area of wind noise. It has kind of been a—the concentration has been on actually doing the stuff that we have been doing for the last 10 years more efficiently. In about, I would say, 1996 or so, we had the expectation that the next-generation of the Cray T90, which would have given us a boost of performance, would have naturally been available, but ended up being canceled.

So it is amazing that we still have a T90 at Ford. That machine is almost 10 years old, and it still does things that, basically, can-

not be done effectively by anything new. So I am looking forward to future advances in this—pushing this type of technology so that we can address the difficult problems that have kind of come to a standstill.

Mr. NEUGEBAUER. As a follow-up to that, because part of the question there is the access to the capability and then developing the capability, is the problem at Ford the capability or the access to existing capability?

Mr. SCARAFINO. The problem is that there is no place we can go to buy something like that. So I—basically, it is really not available, as far as—there isn't something out there that we can't have access to. I mean, we can make arrangements to get access to something, if it is available.

Mr. NEUGEBAUER. As a follow-up to Dr. Reed, what areas or what groups are not able to access existing technologies today because there is not the appropriate infrastructure in place to—for them to be able to access it? Who—what groups would you say?

Dr. REED. Well, we provide access, as I said, essentially to anyone. In fact, one of the hallmarks of the NSF Centers has been that we provide access to researchers regardless of the source of their funding. So researchers who have DOE awards, National Institutes of Health awards, NASA awards, all kinds of people have gained access to the machines over the years via peer review. The thing that is necessary for certain groups to gain access is, if I can use a—perhaps a bad analogy. Here is an old cartoon that shows a monster at the edge of the city with a big sign that says, "You must be at least this tall to attack the city." In order to be able to effectively use high-end computing systems, the application group needs to have sufficient support infrastructure in terms of software development, testing, being able to manage the results, and correlate those with experiments. So it is possible for single investigator researchers to use facilities, but it is much more common for an integrated group of people, perhaps five or ten people, faculty, staff, research associates, to work together to be able to manage a large scale scientific application.

In terms of other components that are necessary, clearly, high-speed networks are the on-ramp for remote access. And as Dr. Freeman noted earlier, one of the things that is increasingly true is that it is not just computing but access to large amounts of data, scientific data archives, from instruments as well as computational data. So the last mile problem, if you will, for universities is part of the issue, having that high-speed capability to be able to move data back and forth, as well as the local capability. The other thing I would say, the most common model of the way machines are used these days is that local researchers typically will have some small copy of a machine on which they do local development and testing. And then they will run it scale at the national center, so that local infrastructure is an additional part of the story.

Mr. NEUGEBAUER. Thank you, Dr. Reed and Mr. Chairman.

Chairman BOEHLERT. Thank you very much.

The gentleman from Oregon, Mr. Wu.

Mr. WU. Thank you very much, Mr. Chairman. And I don't have so much a question right now as a comment. Mr. Chairman, just as you shared a story earlier about the Nobel Prize winner who is,

you know, having access problems and all, I just recall a quite memorable moment. This was probably about 10 years ago in the mid-90's or so. And I was practicing law in the technology field. A friend of mine took me on a tour of—he was a software developer for Intel, and he took me on a tour. And he did supercomputing software. He took me on a tour of their supercomputing inventory, I guess you would call it. I didn't see the production areas, but we were in a room maybe half the size of this one, and it had a lot of HVAC in it and a bunch of black boxes, scaleable black boxes. They were massively parallel supercomputers.

And there was just this one moment when he paused and said— I can't imitate his accent, but he said something to the effect of, "You know, David, you are in the presence of one of humanity's great achievements, like the pyramids." And I will just always—I will remember that. It was striking to me, because Intel got out of the business a little while later. And the concern, I guess, besides the recollection, the concern I want to express is that, you know, civilizations and technologies ebb and flow, and the pyramids are still there. The civilization that created them stopped doing things like that, and for the most part, is no longer around. We had a one shot—well, not a one shot, but a limited series of shots of the moon in the—our space program, there is a sense that there is a high water mark there, and perhaps we have receded from that. I have no sense that our computing industry is necessarily at a high-water mark and receding.

But I just want to encourage you all and this Congress and our government in general to support a sustained push in this and other arenas so that we do not look back at some distant future date and, like me, think back to that supercomputing inventory where the business no longer exists. And you know, people in business have to make business decisions. But as a society, there is a necessary infrastructure where we need—whether it is a space program, whether it is biology, whether it is computer technology, we need a long, sustained push. And it is not necessarily sexy. Those black boxes were not particularly notable from the outside. And I just want to encourage you all and encourage this Congress to sustain that necessary long-term push for the long-term health of our society.

And thank you. I yield back the balance of my time, Mr. Chairman.

Chairman BOEHLERT. The Chair wishes to identify with the distinguished gentleman's remarks. I think the Committee does.

The Chair recognizes the distinguished gentleman from Georgia, Dr. Gingrey.

Dr. GINGREY. Thank you, Mr. Chairman.

Dr. Freeman, I am going to fuss at my staff, because they didn't call to my attention that you were the founding Dean of the College of Computing at the Georgia Institute of Technology, my alma mater.

Dr. FREEMAN. Well, I will do the same. My staff didn't alert me that you were a Georgia Tech graduate. But I discovered that early this morning.

Dr. GINGREY. Had I known that, I would have taken the liberty, as my colleague, Mr. Johnson, did earlier, in introducing Dr. Reed.

I have, actually, two children who are also Georgia Tech graduates. Now I have to admit, when I was at Georgia Tech, it was in the days of the slide rule. And if you wanted to see a teraflop, you just went to a Georgia Tech gym meet.

Dr. FREEMAN. It is probably still true.

Dr. GINGREY. And even today, I have to say that I still think a laptop is a supercomputer. But the staff has done a good job of feeding me a couple of questions. And I want to direct the first one to you, Dr. Freeman.

You stated emphatically that the National Science Foundation is committed to stable, protected funding for NSF-supported Supercomputer Centers. Yet we hear from scientists and from others in the Administration that NSF's focus in the future will be on grid computing and putting high bandwidth connections between the centers and not on supporting top-of-the-line supercomputers at the centers. At what level do you currently fund the Supercomputing Centers? Are you committed to maintaining that level of support separate from the grid computer effort over the next several years? And does the National Science Foundation plan to support future purchases of new supercomputers at the centers?

Dr. FREEMAN. Let me note, Dr. Gingrey, that there are a number of ways in which the Supercomputer Centers, specifically San Diego, Illinois, headed by my colleague here, and at Pittsburgh, there are a number of ways in which they are supported. They have grants and contracts with a number of agencies as well as with a number of different parts of the National Science Foundation. Specifically, the program that is under my peer review has been called the PACI program, the Partnerships for Advanced Computational Infrastructure. That is a program that, for the two partnerships, one located at Illinois under Dr. Reed's direction, one located at San Diego under Dr. Fran Berman's direction, each of those partnerships receives about $35 million a year under the current arrangements.

Some of that money is then reallocated as it were by them to some of those partners, other universities that either operate machines or participate in the development of the advanced software, both applications and the underlying enabling software. In line with our recommendations of our distinguished panel, the Atkins Committee, we are in the process of restructuring the way in which that support is provided. We have, in no way, indicated that the total amount of money going to those activities would decrease. In fact, in speeches both to Dr. Berman and Dr. Reed's partnerships groups, I have indicated our intent to do everything possible to see that the amount of money going into the activities in which they are participating would grow over time.

Dr. GINGREY. And a follow-up question, too, Dr. Freeman. The testimony today has emphasized that high performance computers are not just about hardware and big machines. Dr. Reed and Dr. Orbach both described the importance of investing in research on supercomputer architecture and on the algorithms and software needed to run the computers more effectively. What agencies support this sort of research? Does the industry do or support this sort of research? And is NSF responsible for supporting fundamental computer science research in these areas?

Dr. FREEMAN. The—your last statement is the core characterization of the directorate that I head. Our budget of about $600 million a year, in round numbers, current year. About $400 million of that goes largely to the support of fundamental computer science and computer engineering research in this country. Overall, that supports slightly more than half of the fundamental computer science research in this country.

Let me come back to the opening part of your last question there and emphasize, and I had wanted to in response to one or two earlier questions, I understand that today's hearing is focused on high-end computing, on supercomputers, as it should be. But I believe it is extremely important to keep in mind that there must be a balance, one, and two, that without networking, without software, there is no access. The Japanese supercomputer is, in fact, a case in point. It may be the fastest machine on the planet at the moment, at least in the public domain that we can speak about, but it is not connected to the Internet. So if someone in your district wants to go and use that computer, they would have to travel to Japan.

Secondly, it basically doesn't have much software. And many of the inquiries we have been getting from the scientific community is—are of the form, "How can we put some of the software that we know how to build in this country, how can we put that on that Japanese Earth Simulator machine?" So I would use a simple analogy. A Ferrari is a very fast car, but you wouldn't use the Ferrari to transport oil in or to take honey or to do a lot of other things in. So having just high speed is not the only thing.

Dr. GINGREY. Thank you, Dr. Freeman. And thank you, Mr. Chairman.

Mr. EHLERS. [Presiding.] The gentleman's time is expired.

Next, we are pleased to recognize the gentlewoman from California, Ms. Lofgren.

Ms. LOFGREN. Thank you, Mr. Chairman.

First let me apologize for my delay in coming to this hearing. The Judiciary Committee was marking up the permanent Internet tax moratorium bill, and we successfully passed that, so I think that is good news for the high-tech world. But I do think this hearing is extremely important, and I am glad that the hearing has been held.

Like my colleague from Oregon, Mr. Wu, I remember being at Genentech a number of years ago, and with then Vice President Gore. And there were so many smart people, and they had Ph.Ds in various biological fields. And the Vice President asked them what was the most important technology for their work with the human genome, and they spoke with one voice, "Computers." And really, without an adequate investment in high-end computing, we will see a shortfall in the innovation that we need to have in a whole variety of fields. So this—including fusion, I was glad to hear that the Chairman had identified that.

I have really just two quick questions. The first may be a stretch, but we have, in a bipartisan way, been attempting to remove the MTOP measurement as a restriction for exports of supercomputers. And the question really is, and I don't know who is best suited to answer it, whether that misplaced outdated measurement might be

impairing the private sector development of high-end computing, because it is limiting markets. That is the first question. And the second question, which probably goes to Dr. Freeman, is whether we are sufficiently investing in materials science that will, in the long run, form the basis for next-generation supercomputers. So those are my two questions, and whoever can answer them, I would be delighted to——

Dr. FREEMAN. I am afraid I don't have—I am not sure any of us have much information on the export issue. I was——

Ms. LOFGREN. Okay. Fair enough.

Dr. FREEMAN. I must leave it at that. As to the second—your second question was concerning the investment in materials science. Again, that is not my particular area of expertise. I would note that the nanoscience, nanotechnology initiative of the government that NSF, I believe, is leading and that DOE is heavily engaged in, certainly seems to be pushing the boundaries there. The specific question, I am afraid, I don't have any information on it.

Dr. ORBACH. I could add that the material science area is essential for computation development, especially for the new architectures. And here, DARPA plays a very key role in experimenting with new architectures, new structures that will lead to advanced computing efforts. And most of them are controlled by material science issues. This is an area that all of us are investing in heavily.

Ms. LOFGREN. All right. Thank you very much.

Dr. REED. I could, perhaps, provide just a bit of insight into your MTOP question. It does and it doesn't. One of the things that has been true in parallel computing is that one can export individual components and then often third parties will reassemble them to create machines that are then sold internationally, independently of the original vendors in the U.S. The place where it is, perhaps, more of an issue is in monolithic high-performance machines. The vector computers where one, by very obvious means, can't disassemble them and reassemble them elsewhere. And it does have an impact there, but it is a mixed impact.

Ms. LOFGREN. Thank you. I yield back, Mr. Chairman.

Mr. EHLERS. Well, will you yield, just a moment, to me?

Ms. LOFGREN. Certainly.

Mr. EHLERS. Let me follow up on that. The Japanese are not—who are the leaders in vector machines, don't want to have any such restrictions, is that correct?

Dr. REED. That is correct.

Mr. EHLERS. So doesn't that render our restrictions meaningless?

Dr. REED. It does, in many ways. In fact, those Japanese machines are sold in the markets, yes, where the U.S. vendors cannot sell.

Mr. EHLERS. Yeah.

Ms. LOFGREN. Mr. Chairman, I wonder, I don't want to belabor the point, but whether we might consider having some further deliberations on this MTOP issue in the Science Committee, because I think we have a unique perspective to—it is not really the main point of this hearing. But I think we could help the Congress come to grips with this in a way that is thoughtful if we did so, and I——

Mr. EHLERS. Well, if you will yield.

Ms. LOFGREN. Yes.

Mr. EHLERS. Obviously, I don't control the agenda of this——

Ms. LOFGREN. Right.

Mr. EHLERS [continuing]. Committee, but I would say this is just a continuing, ongoing battle. And I think you and I are both engaged in the——

Ms. LOFGREN. Right.

Mr. EHLERS [continuing]. Encryption battle, which—in which the net effect of the controls was to strongly encourage the encryption industry outside of the United States.

Ms. LOFGREN. Correct.

Mr. EHLERS. And it would, frankly, hurt our country. And I don't know—I fought that battle very hard, and I was astounded at how difficult it is to——

Ms. LOFGREN. Right.

Mr. EHLERS [continuing]. Win those battles for people and use national security as the excuse for their suspicions and—or support for their suspicions, and it would be the detriment of our country.

Ms. LOFGREN. I—my time is expired, and I would associate myself with the gentleman's remarks.

Mr. EHLERS. The gentlelady's time is expired.

And next, we turn to Dr. Bartlett, who has returned.

Dr. BARTLETT. Thank you very much. I am sorry I had to leave for an appointment for a few minutes.

In another life about a third of a century ago, I worked for IBM. And we at IBM were concerned that as a company and as a part of this country that we, as IBM in this country, were at risk of losing our superiority in computers to the Japanese. Apparently our worst fears have been realized. And our reasoning for that, and this was a time when we were really premier in the world, was that every year the Japanese were turning out more, and on average, better scientists, mathematicians, and engineers than we were. And we just didn't think that we, at IBM, and we, the United States, could maintain our superiority in computers if this trend continues.

The Japanese now have, as has been noted by our Chairman in his opening statement and by others, the Japanese now have the world's fastest and most efficient computer. How did they do it? Was it people or was it something else?

Dr. FREEMAN. Let me start the responses. First let me note that I believe that while the Japanese may have the fastest computer, that it would be erroneous to say that they have the complete lead in computing, because there are a number of other elements that are extremely important that they do not have the lead in. And indeed, as I noted a moment ago, I believe you were out of the room, in order to utilize that Japanese machine, they are fairly urgent to get some of our American software technology to enable the usage of that machine. I would underscore what Dr. Reed said earlier this morning, and that is that it is not that the Japanese were smarter. Indeed, a number of the engineers, I understand, that worked on the design of that machine, in fact, had spent time in this country as employees. And I know people who know them and don't think they are particularly any geniuses, obviously very good engineers, but that it has been the sustained focused investment that has

been made and that is, if anything, the single key reason why they now have a Ferrari and we may not have.

Dr. BARTLETT. Dr. Orbach.

Dr. ORBACH. I would like to congratulate the Japanese on their achievement. What they did was to build a balanced machine. Instead of going for just the highest speed they could, they recognized the importance of balance, the communication with the memory, the speed of communication between the microprocessors and the memory. They put together a machine that would be suitable for scientific, and actually, I believe, industrial uses. It wasn't just a peak speed machine. It was a beautiful device.

That is what we are trying to achieve ourselves now. We have learned from the Japanese that it can be done. We are experimenting with different architectures, not necessarily the same, to create something balanced, something that would function for problem solution as opposed to just some arbitrary speed criteria.

Dr. REED. I would echo that. The sustained investment is critical. I go back to what Edison said years ago that genius is 1-percent inspiration, 99-percent perspiration. The sustained effort is really what has been critical. And that, I think, is really the lesson for us, that we need to look at the sustained level investment. The broad array of capabilities and computing technology in the U.S., I think it is absolutely true, as Dr. Freeman said, that overall the preponderance of capability in the world remains in the U.S. We have not harnessed it as effectively as we could, and the Japanese is a—machine is a great example of what can happen when the appropriate talent and sustained investment is harnessed.

Mr. SCARAFINO. One of the other items that, I think, had an effect on direction was the way that metrics are used, the way that we compute what we think that the supercomputer is. They really are based on what individual processors can do running specific benchmarks. Microprocessors got very fast. I mean, Moore's Law basically allows them to be faster and faster as there—as the new generations come out. But they don't address how the overall system performs. If you look at the way that the top 500 supercomputer list is, it takes what effectively any one component of a massively parallel machine can do and multiplies it by the number of elements that are in there. And that is a theoretical peak number, but it is something that, in a practical application, can never really be achieved.

And I think that is something they are understanding better now that we have built some of these large machines and now that we have tried to run actual applications on them where we are seeing five percent of peak available. Even with a lot of work in order to get up to those high numbers, we are starting to understand that the metrics that we are using to describe these machines lack some of the characteristics, that they don't really reflect the actual reality of the situation. If you look at the number that the Earth Simulator has with regard to the metrics versus the next fastest machine, I believe the actual reality of the disparity between the two machines is much greater than those numbers show.

Mr. EHLERS. The gentleman's time is expired.

We are pleased to recognize the distinguished gentleman from Tennessee, Mr. Davis.

Mr. DAVIS. Thank you, Mr. Chairman. It is good to be here, and I apologize for being late when most of you gave your testimony, so therefore, I have basically reviewed some of the testimony that has been—that was given while I was at another meeting. It seems like that, here in the Congress, we have four or five meetings at once. I understand the term being at two places at once. Sometimes maybe it is four or five.

Dr. Orbach, I have a couple questions for you. One will deal, obviously, with the supercomputers that—in the different labs, and the one I am most interested, obviously, is the one in Oak Ridge, which is part of the district that I represent. But as I have read your testimony, you mentioned that there is a new sociology that we will need to develop when it comes to supercomputers, a sociology of high-end computation that will probably have to change. What—exactly what do you mean by that when we talk about the new sociology and how we may have to change?

Dr. ORBACH. The size of these computers and their speed is beyond anything that we have really worked with before. And it is my view that it will take teams of scientists to be able to utilize these facilities in their most efficient fashion. And not just the scientists interested in the solution of the particular problem, but applied mathematicians, computer scientists, people who can put together a team to actually utilize these machines. It is my view that these machines will be like light sources or like high energy physics accelerators. Teams will come to the machine and utilize it as users. And by the way, industry, also, will come and use these machines as a user group.

We don't do that now. And when I made earlier reference to opening NERSC, our big machine, up for grand challenge calculations, part of that was directed toward this new sociology. How can we learn how to use these machines most effectively? The amount of data that they can produce is huge. Our ability to understand it in real time is limited, that we need visualization methods to understand what it is we are doing. There is so much coming out. These are things that we really have to get used to learning how to do. And even the Japanese themselves on the Earth Simulator are finding that it is difficult to understand, in real time, what you are doing. And so they are going to visualization methods. It is that that I was referring to.

Mr. DAVIS. Okay. It—when the Earth Simulator first came—obviously most scientists, I think, were aware that Japan was working in this area, somewhat surprised, I think, when it came on line as quickly as it did and it came up as quickly as it did. I know that there are locations throughout America that basically got excited as well in this direction. "We have to be competitive. We have to find or develop a supercomputer or at least the area for it." I—again, I am selfish, so I think our folks, in the area where I am from, at the—which is one of the labs that does a great deal of work. We have got the SNS [Spallation Neutron Source] project there, which is a part of—when you talk about sociology, there are a lot of folks that are visiting scientists. And my hope is that as you start looking at locations, I know there is, what, a $35 or $40 million increase for supercomputers in the budget and the House is looking at the same thing as the Senate. I don't even know what kind of

criteria that you will use as you determine the location: maybe the capabilities, capacity, some of those issues as you look at how you will make that decision of where those dollars will be spent and the lead location of—for our supercomputers.

Dr. ORBACH. We anticipate that there will be more than one lead location. Our investment now is to look at new architectures. The X–1, for example, at Oak Ridge National Laboratory, is but one example. As has been discussed before, the number of vendors in the field is too small. We are encouraging industry to develop interest in high-end computation. As these new vendors come on line and decide to get interested in this, we could see a multi-center development of—and testing of these new facilities. Ultimately, when the big machines are built, and I hope there will be more than one, there will be a competition between the laboratories as to which would be the most effective place for the computation facility to be housed.

Mr. DAVIS. Of course maybe I shouldn't press them further, but we have got cheap electricity—thank you so much.

Do I have more time? I have one more question to Mr. Scarafino.

Mr. SMITH OF MICHIGAN. [Presiding.] You don't. Your time is up, but maybe if it is a short question——

Mr. DAVIS. It would be. Mr. Scarafino, I know that the automobile industry, especially from some of our competitors in Asia, automobiles may be—I don't want to say quieter than the American vehicles, because I drive an American vehicle, probably always will. But are you able to work—I mean, if—do you sense that industry in this country has the connection, the ability that supercomputers are available to you or will be in the future? And would you use this asset in private industry if that was made available?

Mr. SCARAFINO. You made a reference to vehicles being quieter. Is that what——

Mr. DAVIS. I am just saying, when you look at the auto industry in Asia, it seems that their cars—I am getting in trouble with you, aren't I?

Mr. SCARAFINO. No. No. It is a reality and——

Mr. DAVIS. They may be quieter. And some folks think, perhaps——

Mr. SCARAFINO. More durable, at least.

Mr. DAVIS. I don't know about more durable. I mean, I drive—I have always driven American cars, and I don't see that they are any more durable.

Mr. SCARAFINO. Okay.

Mr. DAVIS. I get 250,000 miles. Mine always rode finer before I sold them, so—but it seems that industry in some of the Asian countries that are our competitors are being allowed or given the opportunity to use the assets that may be funded by government more so than private industry. Is that the case, and should that be changed?

Mr. SCARAFINO. Well, I don't know where the funding comes from. I know that some of our challenges, one with regard to the—to noise, for example, has to do with wind noise in particular that we are running on the fastest machines we have access to. And as I had stated before, some of these problems run for two weeks on a computer to get a single answer. Toyota does have a fair number

of vector machines that is—but they are either NEC machines, and they also have Fujitsu machines, which are, you know, basically Japanese vector high-end processors. I don't know exactly what they run on them, so I mean—so it could be implied that they are utilizing these machines in order to create the quieter vehicles they have. But that is—it is kind of speculation on my part.

Mr. SMITH OF MICHIGAN. The Chairman would call on the Chairman of the Subcommittee on Space, Mr. Rohrabacher.

Mr. ROHRABACHER. Thank you very much.

Maybe you can answer a question for me. I have just got some very fundamental questions about computers. Way back when, you know, like when I was a kid, I seem to remember—and that was a long time ago. This is like after World War II. Weren't there computers that—didn't they have cards that you punched out, and was it an "x" and a "y" or something like that or a "y" and a "0"? Or what was it? Do any of you remember what I am talking about? Is that what it is still based on? Does it still go back to that fundamental principle whether two things being punched out on a card and we have built from there?

Dr. ORBACH. Well, the—as I am older than you are, and I, believe me, carried around those crates of punched cards from one machine to another. They—if it was punched, it was a "1". If it wasn't punched, it was a "0". That digital——

Mr. ROHRABACHER. Okay. "1" and "0".

Dr. ORBACH. That digital structure still is underpinning our computational methods.

Mr. ROHRABACHER. All—so everything we are talking about still goes back to that "1" and that "0"?

Dr. ORBACH. That "1" and that "0". Now there are quantum computation methods that are being explored that will bring an additional richness to that. These machines—NSF is supporting research in that area. DARPA especially is supporting research on some of these advanced concepts that will carry us beyond the "0" to "1", the binary methods that we have been using over the years. And the richness of some of these methods is very exciting. It will be interesting to see how they develop.

Mr. ROHRABACHER. So we are right now thinking about that basic fundamental may change now and take us into something that we—may even be more spectacular than what we have got?

Dr. ORBACH. Yes.

Mr. ROHRABACHER. Let us think about that. I also remember when I—of course we are talking—I remember people talking about a lot of everything that we were doing in computers was based on sand, and making glass out of sand and electricity being run through sand that was reconfigured. Is that still the case?

Dr. ORBACH. Well, I think they were referring to the silicon. Silicon dioxide is sand.

Mr. ROHRABACHER. Right.

Dr. ORBACH. But yes, the silicon-based machines, microchips are still the base of choice, though there are other semiconductors that are used. But it is very plentiful, but it has to be purified to a degree that makes it tricky and expensive.

Mr. ROHRABACHER. We have turned this "1" and "0" and sand into enormous amounts of work and wealth creation. That is mind

boggling. I used to be a speechwriter for Ronald Reagan back in the 80's. And that is getting afar back. Now I was asked to go to the Soviet Union, back when it was the Soviet Union, right after Gorbachev took over and things were beginning to fall. And something really—I had a little experience there that really taught me—and we, because there had been all of this talk about supercomputers and—at that time. The supercomputer was going to change everything. But when I went over to Russia, I took with me a bottle of peanut butter, because I knew they couldn't make peanut butter. At the right moment, I had a bunch of college kids talking to me about America. I pulled out the jar of peanut butter, and I—you know, you can imagine if you have never tasted peanut butter.

But getting on with the story, one of them came up to me after hovering with the other kids there and said, "What are the black marks on the side of the peanut butter jar?" I said, "Well, that is a bar code. And that is where every time I go to the store and buy a jar of peanut butter at the food store, it itemizes my bill for me. The computer itemizes the bill, and an inventory is notified that there is an item that has been sold. And you know, that makes everything easier." And the kids got together. These are college kids in Russia at the time. "You know, that is why we don't trust you Americans. You know, you are lying to us all of the time. Computers at a food store? Give me a break." And they could not believe that we had—and I went to their food store and of course they were using abacuses and things like that. And it is—you know, that impressed me.

That is when I knew we were going to win the Cold War. I mean, there was no doubt about that. It wasn't just the peanut butter. It was the fact that we had computers being put to use across the board in our economy and food stores where they couldn't even think about doing it at food stores.

Now today, when we are—I am very happy that Mr.—for Dr. Freeman and Dr. Orbach have talked about balance. Are we balanced in the fact that we are working to make sure that there is a widespread benefit to this sand and computer-generated numbers? A widespread benefit versus only putting our eggs and trying to create a pyramid, so to speak?

Dr. ORBACH. If I could respond, I hope we are. We are certainly encouraging industry, as you described it, to join us in this quest for seeing what can be accomplished at these speeds. Both NSF and DOE and actually the Office of Science and Technology Policy have included industry in the development of these new machines for exactly the purposes that you recognized before. And as you have heard from the representative from Ford, but also from GE and GM and other companies, there is widespread recognition that these machines will give us economic competitiveness, exactly as you described, that nobody else can get.

Mr. ROHRABACHER. Across the board rather than just——

Dr. ORBACH. Across the board.

Mr. ROHRABACHER [continuing]. At the huge level?

Dr. ORBACH. We won't have to build models. Do you remember the wind tunnels with airplanes and they put them in? If we have enough speed, we can model that computationally and do it in a

matter of hours or weeks, at the most, instead of years. We can save huge amounts of monies on these, what we call, virtual prototypes that will free us from having to build specific models.

Mr. ROHRABACHER. And then we could spend that money elsewhere on something else?

Mr. SMITH OF MICHIGAN. Yeah, I was going to ask Mr. Scarafino if—how many crash dummies the Cray computer has saved, but you—are you——

Mr. ROHRABACHER. I guess I am done. Let me just note that Tim Johnson wanted me to, for sure, say that—to Professor Reed that your institution has such a great record of achievement and wanted me to ask one question that is—and this is the easiest question you will ever get from Dana Rohrabacher, just what do you attribute your great record to?

Dr. REED. You are right. That is an easy question. Well, let me hop back to give you a serious answer, because I think it is important. I mean, the success of the high-end computing infrastructure has really rested on a combination of the fact that there is a high-end infrastructure but also on the critical mass of people. You bring great people together to work with the world class infrastructure and on really critical applications, and interesting things happen. And that is where the web browser spun out of NCSA. It is where a lot of the technologies that will enable cyberinfrastructure have spun out, things like large-scale data mining. There are lots of collaborations with major industrial partners that, as Dr. Orbach said, it saved major corporations lots of money by allowing them to think smarter, to avoid killing crash test dummies, by applying technology in creative ways. And so what we look for are opportunities where the combination of people and technology can apply.

You mentioned the broad spread applicability of computing. The next big wave beyond, sort of, the grocery store and personal computer kind of thing is—computing is really becoming ubiquitous. One question I often ask people is: "How many computers do you own?" And if they have kids, the answer might be three or four. But if you bought a car in the last 20 years, you have anti-lock brakes, you have an electronic thermostat. The answer is really hundreds. And that proliferation of technology, the ultimate success is the extent to which it is invisible and it enriches and empowers people's lives. And that is the broad tie that really is where things like NCSA and other centers have had an impact.

Mr. SMITH OF MICHIGAN. Gentlemen, you have been very patient today. Thank you for your time and patience. You only have one more questioner to face.

Let me—I want to get into the competitiveness and the economic stimulus to the country that might have this computing. Let me start out with the Earth Simulator. Do—does Japan sell time on that to our scientists in this country?

Dr. ORBACH. The Japanese have made time available. Mr. Sato, who is the director of that machine, has about 16 to 20 percent of the machine available to him. And he has been very—he spent time in the United States, in fact, studying science, computer science here. He has made the machine available to our scientists gratis. The somewhat worrisome feature is that our scientists have to go to Yokohama to use it, and the discoveries are going to be

made there. That is what we are worried about. One of the interesting structures, again, the sociology of how the machine is being used, is that there are a number of Japanese young scientists who are there while our scientists are developing their codes and using that machine. So they get the advantage of discovery on the site. And as you have already heard, it can't be networked. They did that on purpose. So they have the availability of discovery, but we don't. And——

Mr. SMITH OF MICHIGAN. Well, at least a little more access and a little more convenience. Dr. Reed, somebody, help me understand a little bit. Our—and Dr. Orbach, it is sort of tied in with—I mean, with—has IBM and Argonne and certainly the Berkeley National Lab developed new computer systems that are going to be even faster than the Japanese current model? And with the so-called development of the—a Blue Gene computer, and I—where we go with that computing capability, which is much faster, how does that tie in terms of our ability? Is it a concern of—the first one, I guess the estimated development time is '05 some time and maybe later for the other one. How does that tie in to the needs of this country and the competitive position in terms of high-end computers?

Dr. ORBACH. It is one of the architectures that I mentioned before that we are exploring. We are looking at different architectures to see how efficient they are and how useful they are for particular scientific problems. And indeed, the Blue Gene, with the collaboration you have made reference to, is a very interesting and important component. It is one of four or five different architectural structures that we are looking at.

Dr. REED. It is certainly the case that the Supercomputing Centers that NSF funds have pretty deep collaborations and connections with the DOE efforts. Argonne National Lab that you mentioned is in Illinois. And we have a jointly funded effort with machines deployed at Argonne and at the University of Illinois. And we are also involved in collaborations with them and with NERSC looking at applications of that machine and what architectural features will be met—best matched to allocations. We certainly hope that in the coming years we will have the resources to be able to deploy one of those machines.

Mr. SMITH OF MICHIGAN. Well, it—as I think you know, I chair the Research Subcommittee, that has oversight over the NSF. And Dr. Orbach, as—if it continues to develop that Energy plays a more prominent role in the development of high-end computers, NSF has been very organized to allow a lot of use in our labs. Is Energy going to be in a position? You mentioned 10 percent of potential usage now, but it is going to take not only—it is going to take much more than that if you are going to accommodate the demand for that kind of time. And it is going to take both organization and facilities, I would think.

Dr. ORBACH. I think you are exactly right. We will—we depend on the National Science Foundation network for coupling to our large-scale computers. And we work very closely with them. Our initiative in high-end computation will be open to everybody. Anyone who can—who wishes can compete for time, subject to peer review, on these new high-end machines. And so we will—we are working hand in glove with NSF to make them accessible. The

challenge now is to determine which architecture we want to invest in and create the structure that would enable our scientists to perform their needed calculations.

Mr. SMITH OF MICHIGAN. In—like our basic research that is oft times sophisticated into application in other countries quicker than we do it in this country with our mandate for publishing any time federal money goes into it, likewise give me your impression as we face a more competitive economic market throughout the world with everybody else trying to produce our cars and everything else as efficiently as we do and as quality conscious as we have been. With our allowance of scientists from other countries to use our—whatever computers we develop, how are we going to still have the edge in terms of the application in this country of the greater efficiencies and development of new products and better ways to produce those products? How are we going to make sure that our investment has the advantage in this country?

Dr. FREEMAN. Well, let me just respond. I think certainly in the general case, open science is still, by far, the best. The competitiveness that the United States—the competitive edge that the United States has, in many, many cases, is due specifically to very talented, very bright scientists and engineers who have chosen to come to this country to study at our universities and oftentimes stay so that open science——

Mr. SMITH OF MICHIGAN. Well, except in—that is changing very rapidly. They oftentimes choose to stay, but it is now a lesser option since 9/11, and so that should concern us.

Dr. FREEMAN. In some—I quite agree that that is a concern, but as I noted in my opening remarks, the synergy that has existed between industry and basic research, for example, has been very productive. And we certainly want to see that continue.

Mr. SMITH OF MICHIGAN. Can I get your reaction from Ford Motor Company, Mr. Scarafino?

Mr. SCARAFINO. I am in agreement. What we really need is access to the kind of computing capability that we can give to our engineers. And by giving them those kind of tools, I am very comfortable that they will know exactly how to make the best use of them. And we would remain—provide us with a way of being competitive, even if those tools are available other—in other places in this world.

Mr. SMITH OF MICHIGAN. Well, I did—let me wrap this up by saying, Dr. Freeman, it still concerns me when you said students from other countries come in and often stay here. Right now half of our research through NSF, for example, is done by foreign students. As we try to do what we can to stimulate interest and stimulation for American students to pursue math and science careers, the question earlier was asked who is—why is Japan developing these kind of high-end supercomputers instead of the United States. Can we expect in the future to see this kind of development, both of hardware and software, go to countries that give greater emphasis to math and science, whether it is China or whether it is India or some other country that tends to encourage and push students in that direction? Is—do you see that as the trend, Dr. Freeman?

Dr. FREEMAN. Well, yes. If I may respond there, I certainly would agree with you that we need to get more Americans into science and engineering. And as you well know, that is one of the primary emphases of NSF across all fields. The computing field is also certainly one of those where we are doing everything we can to encourage more American students, not to exclude the foreigners that choose to come here, because indeed we need more people in general in the computing arena, but we certainly must get more American citizens into the pipeline.

Mr. SMITH OF MICHIGAN. Let us wrap this up. I—maybe if each one of you have a comment of advice for the Science Committee and for Congress in general in your arena, maybe if you have between 45 seconds or so just to—any last thoughts that you would like to pass on for the record for Congress and the Science Committee.

Dr. ORBACH. First, I would like to thank this committee very much for holding this hearing. I think you have brought out the key issues that we all are concerned with and we are very grateful for having this hearing. I would like to comment that on the competitiveness, American industry has always stepped up to the plate. And they are a partner in the development of high-end computation. And I think that is a very special American trait. If we can keep our industry a partner as we develop these new machines, I think it will show in the marketplace.

Dr. FREEMAN. I would stress three things in closing: supercomputing is important; but secondly, it must be looked at and understood in the broader context that I believe all of us have addressed this morning of storage, networking, et cetera; and third, a very important topic that I believe has only been brought up in the last few moments is that of education. If we do not have the trained people, let alone to create such capabilities, but to utilize such capabilities to make sure that they are applied to, whether it is industry or the most advanced most basic research, if we do not have those people, then it makes no difference what type of supercomputers we have.

Mr. SMITH OF MICHIGAN. Dr. Reed.

Dr. REED. I think we are on the cusp of something truly amazing in what we can do with computational capabilities, some very fundamental questions that are as old as mankind are really close to being within our grasp. I think in order to capitalize on that, we have to look at the level of sustained commitment to build the kind of machines that will tackle those kinds of problems. I think that is going to require coordinated investment and activity R&D across the agencies. And I urge that, you know, we look carefully at how to make sure that happens so we can capitalize on the opportunity and maintain the kind of competitive edge that we have historically had.

Mr. SCARAFINO. I would like to thank you for having this hearing. It is good to know that the country is understanding the importance of this area and is willing to basically make more progress in it.

Mr. SMITH OF MICHIGAN. Gentlemen, again, thank you for your time and consideration and your patience. With that, the Committee is adjourned.

[Whereupon, at 12:31 p.m., the Committee was adjourned.]

Appendix 1:

ANSWERS TO POST-HEARING QUESTIONS

ANSWERS TO POST-HEARING QUESTIONS

Responses by Raymond L. Orbach, Director, Office of Science, Department of Energy

Questions submitted by Chairman Sherwood Boehlert

Q1a. At the hearing you spoke of developing and purchasing new supercomputers to be installed at Department of Energy (DOE) labs and of making these computers broadly available to U.S. researchers. When would such computers be installed and be open for general use?

A1a. Our IBM system at NERSC at Lawrence Berkeley National Laboratory is currently open for general use—and has recently been upgraded to over 6,000 processors, making it one of the largest machines available for open science. In fact, we are setting aside 10 percent of this resource for large problems with the potential for high scientific impact. All researchers, regardless of their source of funding may apply.

We are currently evaluating a small Cray X1 system at the Center for Computational Sciences at Oak Ridge National Laboratory. As we transition from evaluation and research into broader use, this system will also be available. The initial evaluation is being done according to an open plan and involves many researchers from the general community already. Given the availability of additional funding, more capable systems could be made accessible as early as FY 2005.

The Pacific Northwest National Laboratory recently announced that a new Hewlett-Packard supercomputer with nearly 2,000 processors has been installed in the Environmental Molecular Sciences Laboratory (EMSL) and is now available to users. As a National Scientific User Facility, the resources within the EMSL are available to the general scientific community to conduct research in the environmental molecular sciences and other significant areas.

Q1b. On these machines, would a certain percentage of the time be set aside for scientists associated with DOE? What priorities would determine who received what amount of time? What peer review mechanisms would be used to make awards?

A1b. A percentage of time on these machines would be set aside for scientists associated with DOE. We have procedures in place to allocate time on these resources—to both DOE and non-DOE scientists. The process is described below.

We have an Office of Science allocation plan in place. It allocates time to our Associate Directors, who in turn allocate it to researchers who are working on their science programs.

We have a peer review mechanism currently in place at NERSC. It has withstood the tests of time and we plan to continue to use it as long as it serves the community well.

All Principal Investigators funded by the Office of Science are eligible to apply for an allocation of NERSC resources. In addition, researchers who aren't directly supported by DOE SC but whose projects are relevant to the mission of the Office of Science may apply to use NERSC.

Four types of awards will be made in FY 2004.

1. Innovative and Novel Computational Impact on Theory and Experiment—INCITE Awards:

Ten percent of the NERSC resources have been reserved for a new Office of Science program entitled *Innovative and Novel Computational Impact on Theory and Experiment* (INCITE), which will award a total of 4.5 million processor hours and 100 terabytes of mass storage on the systems described at *http://www.nersc.gov/*. The program seeks computationally intensive research projects of large scale, with no requirement of current Department of Energy sponsorship, that can make high-impact scientific advances through the use of a large allocation of computer time and data storage at the NERSC facility. A small number of large awards is anticipated.

Successful proposals will describe high-impact scientific research in terms suitable for peer review in the area of research and also appropriate for general scientific review comparing them with proposals in other disciplines. Applicants must also present evidence that they can make effective use of a major fraction of the 6,656 processors of the main high performance computing facility at NERSC. Applicant codes must be demonstrably ready to run in a massively parallel manner on that computer (an *IBM* system).

Principal investigators engaged in scientific research with the intent to publish results in the open peer-reviewed literature are eligible. This program specifically encourages proposals from universities and other research institutions without any requirement of current sponsorship by the Office of Science of the Department of Energy, which sponsors the NERSC Center.

2. DOE Base Awards:

Sixty percent of the NERSC resources that are allocated by DOE go to projects in the DOE Office of Science Base Research Program. DOE Base awards are made by the Office of Science Program Managers and DOE's *Supercomputing Allocations Committee.*

The largest production awards are called **Class A** awards. These projects are each awarded three percent or more of NERSC resources; collectively they receive about 50 percent of the resources. In addition, they may receive extra support from NERSC, such as special visualization or consulting services.

Class B DOE Base projects are awarded between 0.1 percent and three percent of NERSC resources.

All DOE Base requests are reviewed by DOE's *Computational Review Panel* (CORP), which consists of computational scientists, computer scientists, applied mathematicians and NERSC staff. The CORP provides computational ratings to the DOE Program Managers on the computational approach, optimization, scalability, and communications characteristics of their codes. They rate how well the code description questions have been answered.

Projects are rated on a scale of one to five.

3. SciDAC Awards:

Twenty percent of the NERSC resources that are allocated by DOE go to Scientific Discovery through Advanced Computing (SciDAC) projects, which are a set of coordinated investments across all Office of Science mission areas with the goal of achieving breakthrough scientific advances via computer simulation that were impossible using theoretical or laboratory studies alone. SciDAC awards are made by the SciDAC Resource Allocation Management Team. SciDAC requests are not reviewed by the CORP.

4. Startup Awards:

Less than one percent of the NERSC resources are awarded to principal investigators who wish to investigate using NERSC resources for new projects. For FY 2004, the maximum startup awards are 20,000 processor hours and 5,000 Storage Resource Units. A request for a startup repository can be made at any time during the year; decisions for startup requests are made within a month of submission by NERSC staff and can last up to 18 months.

Q1c. Before the new supercomputers are installed, which DOE facilities (and what percentage of time on these facilities) will be available to researchers not receiving DOE grants or not working on problems directly related to DOE missions?

A1c. Currently, 10 percent of the time at NERSC is open to all researchers. The Office of Science program offices are free to go beyond that number when they allocate their portions of the resources: NERSC is currently the only production computing facility funded by ASCR. Should additional production resources become available, they will be allocated in a similar fashion. Time will also be made available on our evaluation machines, in a manner consistent with the accomplishment of the evaluation work.

Because the new Hewlett-Packard supercomputer at PNNL is located within the EMSL, and the EMSL is a National Scientific User Facility, the new supercomputer is available to the general scientific community. Although non-DOE funded investigators may apply for time on the new system, applications must be relevant to the environmental problems and research needs of DOE and the Nation.

Q1d. Are you working with the National Science Foundation to ensure that scientists continually have access to a full range of high-performance computing capabilities?

A1d. Yes. We have formal interactions with NSF (and other agencies) through the National Coordination Office under the auspices of OSTP, and numerous informal interactions in our planning and review process.

Q2. What is the level of demand for cycle time on all of the DOE Office of Science supercomputers? Which scientists and applications are users of the most time?

How will you continue to support the DOE scientific communities and priorities while opening up machines to more general use?

A2. The level of demand for cycles on Office of Science computers is growing exponentially—increasing by an order of magnitude about every 2.5 to 3 years. We are currently operating at full capacity, so managing our out-year demand for additional cycles will be challenging.

Our biggest users include those scientists engaged in research on:

Accelerator design for high energy physics;

Quantum chromodynamics;

Fusion energy;

Climate simulations;

Supernova simulations;

Materials science; and

Nuclear physics.

Scientists in these research areas often use hundreds—and occasionally use thousands—of processors and use weeks to months of resource time.

Continuing to support the DOE scientific communities and priorities while opening up machines to more general use will be a challenge. We are committed to staying this course—because our open science research depends on it. We will ensure that the resources are put to those uses which uniquely exploit them—and do the most to advance the frontiers of science.

ANSWERS TO POST-HEARING QUESTIONS

Responses by Peter A. Freeman, Assistant Director, Computer and Information Science and Engineering Directorate, National Science Foundation

Questions submitted by Chairman Sherwood Boehlert

Q1. Witnesses at the hearing and staff at the Office of Science and Technology Policy have indicated that the National Science Foundation (NSF) may be underfunding research on new architectures for high-performance computing. How much is NSF spending in this area? How did you determine this spending level? How does this level of funding relate to your assessment of the need for research in this area? Are any other agencies or companies funding this sort of research? To what extent are NSF's programs coordinated with these other activities? If NSF does not invest more in this area, are other agencies or private entities likely to fill the gap?

A1. NSF funding impacting new architectures for high-performance computing can be tallied in several ways. In the narrowest sense, funding that may lead most directly to new computer processors is estimated to be around $5 million in FY 2003. Funding at this level is provided from the Computer Systems Architecture program.

In addition to specific research on computer systems architectures, NSF supports research and education in other areas critical to progress in high-end computing, including algorithms, software, and systems. Algorithmic research is essential to enable the transformation of traditional sequential algorithms into parallel ones, and to find algorithms for new classes of problems and new types of architectures. Research on compilers, operating systems, networking, software development environments, and system tools, is also essential if high-end computing is to succeed. It is difficult or impossible to separate out which of this research is directly applicable to high-end computing, but we estimate it to be at least $50M in FY03.

In the broadest sense, NSF funding of the Extensible Terascale Facility can be viewed as research on a specific architecture for high-performance computing. As detailed below,

$126M has been spent on this advanced R&D project to date. The high level objective is to show the feasibility of providing high-performance computing capability via a highly-interconnected, distributed set of computational resources.

In addition to this well-defined project, NSF invests in higher-risk, longer-term research to ensure that innovation is possible years from now. Examples showing promise for high performance computing applications include nanoscale science and engineering, quantum computing, and bio-inspired computing. NSF investments in these areas are in excess of $100 million.

As with all NSF investments, funding levels for computer systems architecture-related activities have been and continue to be determined by a combination of inputs from the communities using high-performance computing in their research, communities with expertise in developing architectures and their supporting technologies, interactions with other agencies, inputs from Congress, and the funds that are available.

NSF's investments in high-end-computing-related research complement investments in computer system architectures, algorithms, parallel software, etc. being made by other federal agencies, especially DARPA, DOD, DOE, and by industry. As in other fields of science and engineering, NSF draws upon partnership efforts to leverage these investments. Specifically in the area of high-end computing they are coordinated through the High-End Computing program component area of the Networking and Information Technology Research and Development (NITRD) Working Group. This partnership has recently been strengthened through the focused activities of the High-End Computing Revitalization Task Force.

As part of the NITRD interagency coordination effort, NSF is considered to have over $200M in "high-end computing infrastructure and applications" and nearly $100M in "high-end computing research and development."

Finally, NSF capitalizes upon the outcomes of Federal Government-supported nearer-term high-end computing research and development activities enabled by its sister agencies, in the support and deployment of high-end computing systems like those currently provided by the National Center for Supercomputing Applications (NCSA), the San Diego Supercomputing Center (SDSC), the Pittsburgh Supercomputing Center (PSC) and the National Center for Atmospheric Research (NCAR) to meet the research and education user needs of the broad science and engineering community.

Historically, mission-oriented agencies such as DOD and DOE have driven investment in the development of operational high-end architectures, with NSF providing the longer-term, higher-risk investments in the basic research and education that must support these operational developments. Industry has responded to the development of new architectures primarily in response to markets, either commercial or governmental. This balanced approach, with swings back and forth between more or less investment in the long-term, has worked well in the past. With appropriate adjustments in response to validated needs, we believe it will continue to work well in the future.

Q2. Different scientific applications need different types of supercomputers, and scientists from many fields use the supercomputing capabilities supported by NSF within CISE. What role does input from these different scientific communities play in the decisions about what types of high-performance computing capabilities are supported by CISE?

A2. As the question indicates, different scientific applications need and prefer different types of supercomputing capabilities. NSF's support for high performance computing in service of science and engineering research and education has been driven by the diverse scientific opportunities and applications of the user community. Informed by user community needs, NSF supports a wide range of high performance computing system architectures.[1] The centers and other facilities supported by NSF are close to the end-users and thus the decisions as to what capabilities to provide are being made as close as possible to the community.

NSF continues to be the largest provider of access to a diverse set of supercomputing platforms for open academic science and engineering research and education in the U.S. In FY 2002, NSF provided over 152 million normalized service units to over 3,000 users involved in 1200 active awards. NSF's user community includes large numbers of NIH-, NASA- and DOE-funded scientists and engineers.

Examples of the range of computing platforms provided through NSF support are described below:

1. Many NSF users prefer shared memory architecture systems with from 32 to 128 processors with large amounts of memory per processor, since their research codes are not scalable to several thousand processors. That is why NSF provides access to a significant number of large memory 32-way IBM Power 4 systems, all of which are continuously over-subscribed. A new 128 way SMP HP Marvel system is also being installed at PSC.
2. At present, there are on the order of 30 research groups that use NSF-supported supercomputers and have developed highly scalable codes capable of efficiently using systems comprised of thousands of processors, such as the NSF-supported Terascale Computing System (TCS) at PSC. High allocation priority is given to researchers who have developed such research codes. In FY 2002, for example, 5 users accounted for 61 percent of the usage of TCS, for projects in particle physics, materials research, biomolecular simulations and cosmology.
3. Driven by a user community need to run very long jobs requiring large numbers of processors, NCSA, with NSF funding, will deploy an Intel Xeon-based Linux cluster with a peak performance of 17.7 teraflops. This 2,900 processor Linux-based system will be dedicated to users who need large numbers of processors for simulations that may require up to weeks of dedicated time.
4. Another new system, driven again by user needs, is the SDSC DataStar, a 7.5 teraflop/s IBM Regatta system that will be installed this fall. This system will leverage SDSC's leadership in data and knowledge systems to address the growing importance of large-scale data in scientific computing. The new system will be designed to flexibly handle both data-intensive and traditional compute-intensive applications. SDSC's Storage Area Network, or SAN, will provide 500 terabytes of online disk, and six petabytes of archival storage.
5. As one of the world's first demonstrations of a distributed, heterogeneous grid computing system, the NSF-supported Extensible Terascale Facility (ETF) will provide access to over 20 teraflops of computing capability by the end of FY 2004. ETF provides a distributed grid-enabled environment, with the largest single compute cluster being a 10 teraflop IA–64 Madison cluster at NCSA. This system will provide an integrated environment that provides

[1] Experience has shown that the needs of the community are quite diverse, and cannot effectively be met by providing only a single type of system.

unparalleled scientific opportunity to the science and engineering community.

Q3. What is the level of demand for cycle time on each of the NSF-supported supercomputers? Which scientists and applications are users of the most time? Is the demand growing? Do you have a plan to provide the capabilities to meet existing and future levels of demand?

A3. Table 1 below describes current demand based on the number of CPU hours available, the number requested and the number of CPU hours allocated during FY 2003. The resources are allocated by a National Resource Allocation Committee (NRAC) that meets twice per year. Table 1 contains allocations for all CISE-supported systems, some of which are located at SDSC, NCSA and PSC, others of which are located at other partner sites. In general the ratio of requested to allocated CPU hours ranges from 1.4–2.1. Additionally, while it is not a PACI site, the National Center for Atmospheric Research (NCAR) supercomputer has a Top 500 ranking of 13.[2]

Table 1

Service Unit (CPU Hours) Allocations for PACI Computer Systems in FY 2003

System	Top 500 Ranking *(as of 9/03)*	CPU Hours Available	CPU Hours Requested	CPU Allocated Allocated
PSC-TCS1	9	22,441,500	35,058,614	22,370,777
PSC-T3E	360	982,650	772,650	1,002,650
TX-Power4	Not Rated	200,000	968,400	195,000
SDSC Blue Horizon	62	5,000,000	14,987,630	5,048,676
WI Condor	Not Rated	950,000	346,480	996,480
KY-Superdome	238	176,500	58,700	258,700
NCSA-IA32	135	6,155,000	5,646,500	6,175,100
NCSA-IA64	111	2,005,500	7,918,703	1,993,960
MI-SP3 (Power3)	Not Rated	100,000	136,000	125,000
MI-SP2 (Power2)	Not Rated	800,000	300,100	210,000
MI AMD Linux	163	1,900,000	288,000	1,398,000
NCSA-p690	99	1,840,000	6,059,262	1,868,000
BU P690	451	125,000	25,000	75,000
BU-SP	Not Rated	75,000	245,082	75,000

Table 2 overleaf provides a ranked listing of the top 25 users who received allocations on these systems[3] in FY 2002. Their utilization is quoted in normalized service units, taking into account the differential capabilities of the systems. Each investigator generally has accounts on multiple systems. They cover a broad range of disciplines including: particle physics, protein biomolecular simulations, severe storm prediction, and engineering. Many of the users of NSF-supported supercomputing centers receive research funding from NIH and DOE. Many of the DOE-funded researchers also obtain significant allocations through NERSC at LLBL.

---

[2] http://www.top500.org/dlist/2003/06/

[3] These resources are under-subscribed by the geosciences research community since the majority of that community uses NSF-supported NCAR supercomputing capabilities.

Table 2

Normalized Units used by Top 25 Users in FY 2002

Rank	Principal Investigator (PI) Name	Area of Research	Total Usage Normalized Service Units
1	Sugar, Robert L.	QCD, Subatomic Physics	30,586,597
2	Schulten, Klaus J.	Biological Simulations	13,011,332
3	Klein, Michael L.	Materials Research	4,941,263
4	Duan, Yong	Biological Simulations	4,706,803
5	Norman, Michael	Cosmology	3,556,473
6	Scheraga, Harold A.	Protein Folding	3,456,817
7	Toomre, Juri	Solar Atmospheres	2,823,629
8	Joannopoulos, John D.	Materials Research	2,774,981
9	Liu, Keh-Fei	QCD, Subatomic Physics	2,119,550
10	Moser, Robert	Turbulence, Applied Mechanics	1,841,523
11	Maier, Thomas	Materials Research	1,753,715
12	Kinoshita, Toichiro	Subatomic Physics	1,749,045
13	Campbell, David K.	Materials Research	1,734,862
14	Tafti, Danesh	Mechanical Engineering	1,655,110
15	Droegemeier, Kelvin K.	Severe Storm Prediction	1,585,282
16	Kogut, John	QCD, Subatomic Physics	1,510,741
17	Cheatham, Thomas E.	Biological Simulatons	1,267,335
18	Seidel, Harry Edward	Cosomology	1,257,989
19	Karniadakis, George E.	Mechanical Engineering	1,158,103
20	Jakobsson, Eric G.	Biological Simulations	1,052,717
21	Hu, Chi-Yu	Nuclear Particle Physics	985,547
22	Suen, Wai-Mo	Cosmology	898,458
23	Hu, Chi Yu	Nuclear Particle Physics	894,302
24	Schlick, Tamar	Biological Simulations	886,818
25	Roux, Benoit	Biological Simulations	867,576

Demand for supercomputing resources continues to grow each year. It has increased from 20 million CPU hours in FY 1999 to over 60 million CPU hours in FY 2002. During that same period of time the capacity within NSF-supported supercomputer centers has increased from 10 million CPU hours in FY 1999 to 45 million in FY 2002. Anticipated capacity in 2003 is 62 million CPU hours. With the installation of the Dell cluster at NCSA in the current FY, the installation of DataStar at SDSC by early next year, and the completion of ETF construction by the end of FY

2004, the combined capacity will grow by an additional 55 million CPU hours (almost doubling the capacity) by the end of FY 2004.

Q4. Do you plan to provide funding for the purchase of new supercomputers for the supercomputing centers?

A4. NSF's commitment to providing enhanced support for state-of-the-art high-performance computing, in the context of cyberinfrastructure, remains exceedingly strong. Following the trends of the last few decades, the agency anticipates providing strong support for technology refresh and upgrade of existing NSF-supported high performance computing resources for the foreseeable future.

Q5. How is NSF working with the department of Energy (DOE) Office of Science to ensure that scientists have access to a full range of high-performance computing capabilities? What are the short-term and long-term plans to provide this full range of capabilities?

A5. Building on a strong interagency partnership that focuses on high-end computing, NSF is continuing to work closely with DOE to provide high performance computing capabilities to the broad science and engineering community. NSF and DOE staff meet regularly through both formal (like NITRD and the High-End Computing Revitalization Task Force) and informal processes to discuss and coordinate current and future plans.

Demonstrating the strength of this relationship, in the first instantiation of ETF, NSF made an award to four institutions that included Argonne National Laboratory (ANL). The ANL component of ETF supported the construction of a 1.25 teraflop IA–64 system with a 96-processor visualization engine and 20 TB of storage. The award also drew upon essential grid computing expertise available at ANL. In FY 2002 NSF funded the second phase of ETF, in an award that also drew upon expertise at ANL. Rick Stevens from ANL is the current ETF Project Director and Charlie Catlett from ANL is the ETF Project Manager.

Before the end of FY 2003, NSF plans to make an award to Oak Ridge National Laboratory, to interconnect the DOE-funded Spallation Neutron Source and other resources available on the Extensible Terascale Facility. This award, as with all NSF awards, was identified through merit review and will provide unique scientific opportunities for the science and engineering user community.

NSF has already demonstrated a willingness to share NSF-supported computational resources with a number of other federal agencies, including DOE and NIH. For example, the 89 PIs submitting proposals for review to the FY 2003 NRAC allocation meetings cited the following sources of research funding (several PIs had multiple sources of funding): NSF (55), NASA (19), NIH (19), DOE (18), DOD (11), DARPA (4), NOAA (1), NIST (1), and EPA (1). Five of the PIs listed DOE as their sole source of research support, three listed NIH as their sole source of support and two listed NASA as their sole source of support. Thirteen PIs listed no federal sources of support.

Questions submitted by Representative Ken Calvert

Q1. Cyberinfrastructure is by its nature pervasive, complex, long-term, and multi-institutional. The National Science Foundation's (NSF's) cyberinfrastructure effort will require persistence through many machine life-cycles and software evolutions, and across multiple institutions and the staff associated with them.

Q1a. How will the programs that NSF's Computer and Information Science and Engineering (CISE) Directorate creates to accomplish its cyberinfrastructure vision address the challenges associated with complex programs and how will the directorate manage the multi-institutional and long-term collaborative projects?

A1a. Dating back to the 1980's, NSF has sought to match the supply and power of computational services to the demands of the academic scientific and engineering community. The increasing importance of computation and now cyberinfrastructure to scientific discovery, learning and innovation, has guided our strategy to accomplish this. This increasing importance, which has been fueled by advances in computing-communications and information technologies, has led to a number of programmatic changes over the years. For instance, the Partnerships for Advanced Computational Infrastructure (PACI) program was developed to meet an increased demand for computing-related services and assistance to accompany the raw computing power enhancements that technological innovation had provided.

The formulation of a cyberinfrastructure vision represents a new set of opportunities, driven by the needs of the science and engineering community and the capabilities that further technological innovation have provided. NSF will build upon the

scientific and programmatic expertise and capability developed over the past few decades to meet the needs of the largest number of science and engineering researchers and educators.

Promising management approaches, designed to engage multiple institutions in long-term collaborative projects, are currently being discussed in collaboration with the science and engineering community. A number of workshops and town hall meetings, which included representatives from academe, other federal agencies and international organizations, were held over the summer of 2003 to discuss promising approaches.

Q1b. How will NSF and CISE address the need to provide support overtime and institutions so that investments in facilities and expertise can effectively be leveraged?

A1b. As indicated above, promising management approaches, designed to engage multiple institutions in long-term collaborative projects, are currently being discussed in collaboration with the science and engineering community. A number of workshops and town hall meetings, which included representatives from academe, other federal agencies and international organizations, were held over the summer of 2003 to discuss promising approaches. Our goal is to use the most promising approaches identified to ensure that investments in facilities and expertise can be effectively leveraged.

Q2. NSF has already announced that the current Partnership for Advanced Computational Infrastructure (PACI) program will end at the end of fiscal year 2004, and that the new program for TeraGrid partners will begin in fiscal year 2005. However this is less than 15 months away and the supercomputing centers have not heard information about the objectives, structure or recommended budgets for the program.

Q2a. What is the schedule for developing these plans and what can you tell us now about plans for this or other related programs?

A2a. NSF plans to issue detailed guidance to SDSC, NCSA and PSC during the fall of 2003. Guidance will include discussion of means of support to be provided for both ETF, in which SDSC, PSC and NCSA are partners together with other organizations, and for other high performance computing resources and services provided by these centers under NSF support. Some discussions with senior SDSC, PSC and NCSA personnel have already taken place.

Q2b. Will NSF's plans be consistent with the recommendation from the Atkins report, which states in Section 5.3 that "the two existing leading-edge sites (the National Center for Supercomputing Applications and the San Diego Supercomputing Center) and the Pittsburgh Supercomputing Center should continue to be assured of stable, protected funding to provide the highest-end computing resources?

A2b. NSF plans are consistent with the recommendations of the Atkins report in that stable funding for NCSA, SDSC and PSC is anticipated to provide high performance computing capabilities to the science and engineering community.

Q2c. The Atkins report also assumes the centers would operate with an annual budget of approximately $75 million per center. Is this the direction NSF envisions for this program and if so, when is this level of support expected to start and how long will it last?

A2c. NSF is currently developing a five-year planning document for cyberinfrastructure. The development of this plan is informed by the recommendations of the Atkins report, as well as other community input. The first edition of the plan, which will be a living document, should be available in 2004.

Questions submitted by Representative Ralph M. Hall

Q1. The National Science Foundation (NSF) blue ribbon panel on cyberinfrastructure needs, the Atkins Panel, which issued its report this past February, emphasized that NSF has an important role in fostering development and use of high performance computers for broad research and education in the sciences and engineering. The Atkins report strongly recommended that U.S. academic researchers should have access to the most powerful computers at any point in time, rather than 10 times less powerful, "as has often been the case in the last decade," according to the report. To provide the research community with this level of capability, the report recommends funding 5 centers at a level

of $75 million each, of which $50 million would be for hardware upgrades needed to acquire a major new system every 3 or 4 years.

Q1a. What is your response to this recommendation that NSF provide the resources necessary to ensure that high-end computing systems are upgraded with regularity to insure that the research community has access to leading edge computers at all times?

A1a. NSF recognizes that technology refresh is very important in order to keep high-end computing resources and hence related science and engineering research and education, at the state of the art. Over the past five years, NSF has demonstrated its commitment to doing so. Table 3 below provides evidence that this is the case.

Q1b. What level of funding has NSF provided over the past 5 years for upgrading the high-end computer systems at its major computer centers? What is NSF's current plan for support of high-end computer centers for providing hardware upgrades needed to keep them at the leading edge (break out funding by categories for operations and maintenance and hardware upgrades)?

A1b. With funding provided through the PACI program and through the Terascale Initiative, NSF has invested over $210 million on hardware and integrated software upgrades and renewal over the past five years.

Table 3. Hardware Expenditures Over Past 5 Years

PACI & Terascale Equipment Expenditures by Year ($K)						
	FY99	FY00	FY01	FY02	FY03	
PACI Computing	$17,239	$19,183	$16,388	$13,263	$18,214	
Terascale Computing		$36,000	$45,000	$35,000	$10,000	
Total Computing	$17,239	$55,183	$61,388	$48,263	$28,214	
					$210,287	Total Computing

The agency remains committed to technology refresh and upgrades, recognizing that this is essential to realize the promise of the cyberinfrastructure vision.

Q1c. You have announced a reorganization of the Computer and Information Science and Engineering Directorate that combines the computer and network infrastructure divisions. Will support for the operation of state-of-the-art high-end computing facilities continue under this reorganization, or does this signal a change in the priorities? Will funding for FY 2003 and FY 2004 for Partnerships for Advanced Computational Infrastructure (PACT) centers increase, decrease or stay the same under this reorganization relative to FY 2002 funding levels?

A1c. The reorganization of CISE does not signal a change in priorities. It has been proposed in order to support the ever-broadening meaning of "state-of-the-art high-end computing facilities" that the cyberinfrastructure vision illuminates. Support for operations of NSF's state-of-the art high-end computing facilities will continue and grow under this reorganization. The reorganization will provide the CISE directorate with a more focused and integrated organization to deal with the deployed computational, networking, storage and middleware infrastructure essential to cyberinfrastructure.

Q2. You pointed out in your testimony that high-end computing is only one component of deploying an advanced cyberinfrastructure needed for the advancement of science and engineering research and education. Are you satisfied with the priority afforded high-end computing in the current NSF budget that supports cyberinfrastructure development and deployment?

A2. High-end computing is essential to the progress of science and engineering. NSF's budget requests and investments are designed to recognize the crucial role cyberinfrastructure plays in enabling discovery, learning, and innovation across the science and engineering frontier. While need continues to outstrip the available resources, NSF budget requests continue to be very responsive to the needs of the community.

Q3. What do you see NSF's role relative to the Defense Advanced Research Projects Agency and the Department of Energy in supporting research related to the development and use of future U.S. academic research community? What portion of your directorate's budget would you expect to allocate for these purposes?

A3. As in other fields of science and engineering, NSF draws upon partnership efforts to leverage its investments in cyberinfrastructure. For example, the agency's partnership with other federal agencies in the area of high-end computing and networking is nurtured through the High End Computing program component area of the Networking and Information Technology Research and Development (NITRD) Working Group. This partnership has recently been strengthened through the focused activities of the High-End Computing Revitalization Task Force.

These partnership activities strengthen the management and coordination of relevant programs and activities to increase the return on current investments and to maximize the potential of proposed investments. NSF leverages nearer-term high-end computing research and development programs funded by DARPA, DOD and DOE, where government-industry partnerships often create new generations of high-end programming environments, software tools, architectures, and hardware components to realize high-end computing systems that address issues of low efficiency, scalability, software tools and environments, and growing physical constraints. By drawing upon its effective interagency relationships, NSF avoids duplication of effort. NSF focuses its investments in higher-risk, longer-term research investments to ensure that the new innovation is possible years from now. Examples of current high-risk, longer-term basic research showing promise for high performance computing applications include nanoscale science and engineering and bio-inspired computing.

Finally, NSF capitalizes upon the outcomes of Federal Government-supported nearer-term high-end computing research and development activities enabled by its sister agencies, in the support and deployment of high-end computing systems like those currently provided by the National Center for Supercomputing Applications (NCSA), the San Diego Supercomputing Center (SDSC), the Pittsburgh Supercomputing Center (PSC) and the National Center for Atmospheric Research (NCAR) to meet the research and education user needs of the broad science and engineering community.

In preparing its annual budget request, the agency gives careful consideration to funding provided for cyberinfrastructure. As the agency focuses increasing attention on cyberinfrastructure, it is likely that the funds dedicated to the development of an enabling, coherent, coordinated cyberinfrastructure portfolio will grow, in recognition of the importance of cyberinfrastructure to all of science and engineering.

ANSWERS TO POST-HEARING QUESTIONS

Responses by Daniel A. Reed, Director, National Center for Supercomputing Applications, University of Illinois at Urbana-Champaign

Questions submitted by Representative Ralph M. Hall

Q1. You cited in your testimony the National Science Foundation (NSF) blue ribbon panel report that recommended a funding level of $75 million per year to enable a supercomputer center to always have a state-of-the-art computer. What is the annual funding level provided by NSF for your center at present, and how has it varied over time? What has been the funding trend for hardware upgrades?

A1. For FY 2003, we anticipate receiving $34,650,00 for the sixth year of the NSF cooperative agreement for the National Computational Science Alliance (Alliance), one of two Partnerships for Advanced Computational Infrastructure (PACI).[1] I have attached a chart that shows the funding history for NCSA (the leading edge site of the Alliance) and the Alliance from FY 1998 through FY 2004 (which is an estimate based on our program plan for the coming year).

As the chart illustrates, NSF funding for the Alliance was $29 million in FY 1998, $34.3 million in FY 1999, $33.7 million in FY 2000, $35.17 million in FY 2001 and $35.25 million in FY 2002. The annual funding level has remained relatively flat throughout the program's lifetime, and we anticipate the FY 2004 total will be $33.5 million, about $1 million less than FY 2003.

I have also included the budget for hardware during this six-year period. In the initial years of the program, software and maintenance were included in the hardware budget line; this cost was moved to the operations budget in FY 2000. As one can see, the hardware budget includes not only funds for supercomputing (capability computing) hardware, but also support for data storage, networking, desktop support, Access Grid (distributed collaboration) and audio/visual support, and testbed systems.

The support specifically for supercomputing hardware has varied. In FY 2001, supercomputing hardware was funded at $9.08 million, in the current year $11.795 million has been allocated, and the FY 2004 request is for $8.63 million. *Funding for hardware upgrades (i.e., the annual hardware budget) has remained fat, varying between $13 and $15 million.*

It is important to note that when the PACI program began, its participants expected overall NSF support to increase annually during the program's ten-year lifetime (i.e., at a minimum to reflect standard inflation). The original request for FY 1998 from NSF was $34.988 million, and we anticipated it would grow to $49.271 million by FY 2002. That steady growth never materialized, despite substantial increases in human resource costs during that time. Hence, the Alliance and NCSA (and its sister institutions SDSC and NPACI) have experienced *de facto* annual budget cuts.

In addition, NCSA, in collaboration with partners at the San Diego Supercomputing Center, Argonne National Laboratory, and the California Institute of Technology, was chosen by NSF to create the Distributed Terascale Facility or TeraGrid.[2] This three-year award, totaling $53 million, has provided $19,450,500 to NCSA for hardware, storage, networking and software support. This is, however, substantially less than an annual "cost of living" adjustment would have provided in aggregate.

Q2. A reorganization has been announced for the NSF Computer and Information and Engineering Directorate that combines the computer and network infrastructure divisions. Have you seen any evidence of how this reorganization will affect the Partnerships for Advanced Computational Infrastructure (PACI) program, and your center in particular?

A2. The NSF's announcement of the CISE Directorate reorganization is relatively recent. There are not yet enough organizational or funding details for substantive assessment.

NSF has announced that the PACI program will end on September 30, 2004—just over one year from now. We do not know, with any specificity, the nature of NSF's plans for high-end computing in FY 2005 and beyond. The cyberinfrastructure panel, which was charged by NSF with evaluating the centers and outlining future

[1] The cooperative agreement for the other partnership, the National Partnership for Advanced Computational Infrastructure (NPACI) is similar.

[2] The Pittsburgh Supercomputing Center was added in 2002.

directions, recommended long-term, stable funding for the existing centers and an expansion of investments in computing infrastructure. However, to date, we have no specific information on NSF plans, out-year budget requests, or the process by which funding would be secured.

With the announced end of the PACI program only a year away and the details of successor structures not yet known, NCSA, our partner institutions within the Alliance, and those at the San Diego Supercomputer Center and the Pittsburgh Supercomputing Center, remain concerned about the future. We strongly believe NSF must continue to make significant annual investments in high-end computing if the national research community is to have access to the most powerful computational resources. Only with these investments can the research community address breakthrough scientific problems and maintain international leadership.

Q3. You pointed out the importance of deep involvement and coordinated collaboration of computer vendors, national labs and center, and academic researchers, along with multi-agency investments, to develop and deploy the next generation of high-end computer systems. What is your assessment of the current state of coordination in this area, and what recommendations do you have on how to improve the situation?

A3. During the early days of the HPCC program, cross-agency activities were somewhat more coordinated than they are now. Hence, my personal view is that there has not been coordinated collaboration on the development of high-end computing systems across the Federal Government for several years. Some of this compartmentalization is understandable, given the differing missions and goals of the several agencies (Department of Defense, National Security Agency, Department of Energy, National Science Foundation, National Aeronautics and Space Administration, the National Institutes of Health and others) that sponsor research in high-end computing and acquire and utilize high-end computing systems.

As I testified at the hearing, the activities of the current High End Computing Revitalization Task Force (HECRTF) are encouraging. Charting a five-year plan for high-end computing across all relevant federal agencies involved in high-end computing research, development, and applications is critical to the future of high-end computing research, computational science and national security. I was pleased to lead the community input workshop for HECRTF, and I am hopeful the views and suggestions of research community will be incorporated into the vision and recommendations of the Task Force.

However, I believe success rests on more than budgetary coordination. We must also ensure that insights and promising ideas from basic research in high-performance computing are embraced and developed as advanced prototypes. Successful prototypes should then transition to production and procurement. Greater interagency coordination is needed to ensure such transitions occur.

PACI Funding FY1998-FY 2004 ($'s in Thousands)

	FY 1998	FY 1999	FY 2000	FY 2001	FY 2002	FY 2003 Proj.	FY 2004 Req.
Total Cooperative Agreement	**$29,000**	**$34,313**	**$33,704**	**$35,175**	**$35,250**	**$34,650**	**$33,500**
Capital	**$13,000**	**$17,331**	**$18,072**	**$11,506**	**$14,655**	**$15,570**	**$15,000**
Hardware	13,000	13,617	12,813				
Maintenance/Software		4,272	5,213				
Capability Computing	9,020	11,353	10,882	9,076	9,900	7,590	8,630
Data Grid/Storage	1,919	1,583	775	1,050	2,000	5,300	4,078
Access Grid and VR	530	100	200	500	700	700	530
Networking	862	525	456	530	1100	1,300	910
Desktop and AV	669	56	500	350	500	380	302
Testbed Systems				0	455	300	350

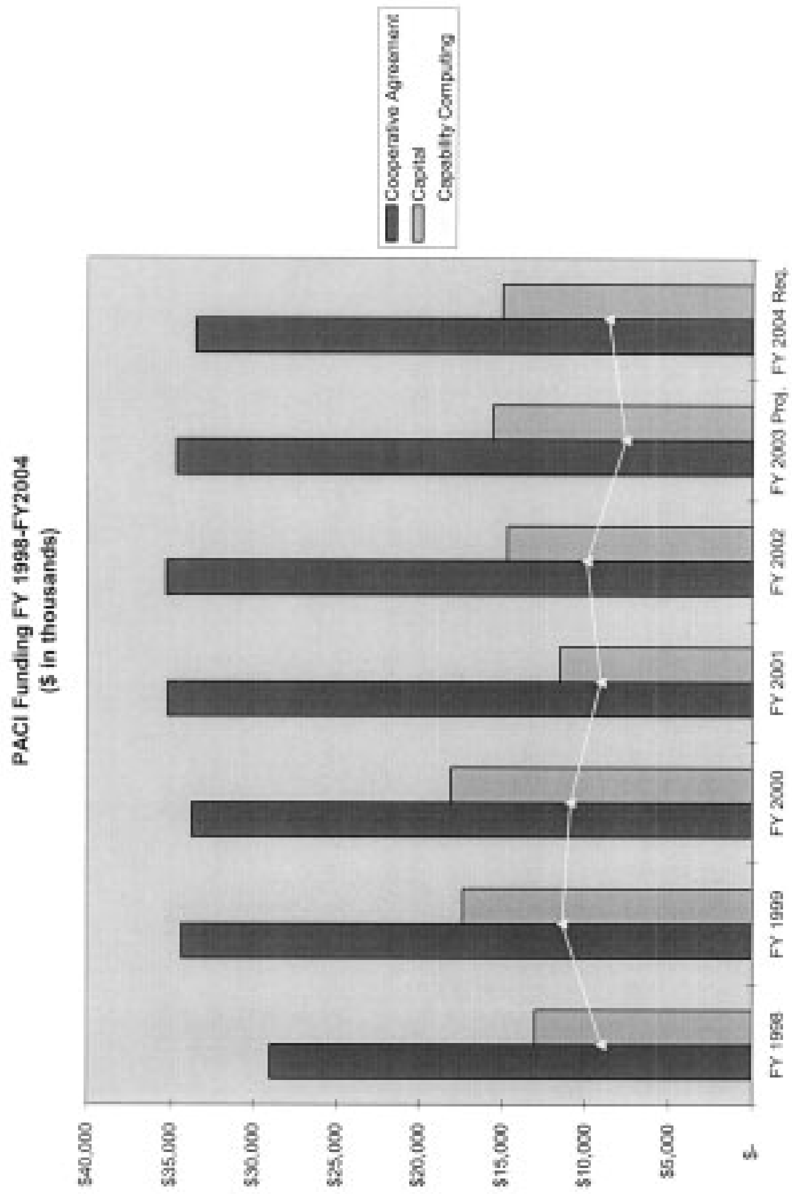
PACI Funding FY 1998-FY2004
($ in thousands)
Cooperative Agreement
Capital
Capability Computing
$40,000
$35,000
$30,000
$25,000
$20,000
$15,000
$10,000
$5,000
$-
FY 1998
FY 1999
FY 2000
FY 2001
FY 2002
FY 2003 Proj.
FY 2004 Req.

Appendix 2:

Additional Material for the Record

ADDITIONAL STATEMENT BY DR. RAYMOND L. ORBACH

There are two aspects to this question: the extent to which supercomputing will aid or accelerate development of nanotechnology (through, e.g., modeling and simulation of nanostructures), and the extent to which nanotechnological advances will contribute to supercomputing in Japan.

With respect to the first issue, there is no question that supercomputers will allow more detailed, complete, and accurate modeling of larger collections of atoms, and permit simulations to cover longer time periods. Both of these capabilities are critical in connecting basic theory to practical experimental data about assembly, structure, and behavior of materials at the nanoscale. There are in fact research projects using Japan's Earth Simulator that address such questions, such as one entitled "Large-scale simulation on the properties of carbon-nanotube" (Kazuo Minami, RIST).

With respect to the impact of Japanese nanotechnology on Japanese supercomputing, the technology appears to not be sufficiently mature for a major impact in the next few years. Nanotechnology products are beginning to head into the marketplace and will affect the computing industry; examples can certainly be found both in Japanese and in American companies. (From Japan, NEC recently announced that they intend to use fuel cells based on a form of carbon nanotube to extend battery cycles in notebook computers from 4 to 40 hours. A recent U.S. example is the development by Motorola of a technology to produce large flat panel displays based on electron emission from carbon nanotubes, for which they are in the process of negotiating licensing agreements.) However, these are currently targeted at the consumer electronics market and may not have immediate impact on supercomputing. For the latter, developments in areas such as heat conduction using nanoscale technologies may have an impact by facilitating cooling of supercomputer architectures.

Beyond the end of the silicon semiconductor "roadmap," in another 10–15 years, the development of molecular electronics may begin to have an impact on all forms of computing. Predictive computer modeling of molecular electronics may well be essential for design and manufacturability of such structures, just as computer-aided design has proven critical to the development of current-day circuits.